BlackDiamond User Guide

Extreme Networks, Inc.
10460 Bandley Drive
Cupertino, California 95014
(888) 257-3000
http://www.extremenetworks.com

Published: September 1998
Part number: 101000-00 Rev. A

Copyright © **Extreme Networks, Inc., 1998.** All rights reserved. No part of this documentation may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without permission from Extreme Networks, Inc.

Extreme Networks, ExtremeWare, BlackDiamond, Summit, SummitLink, ExtremeWare Vista, Summit Virtual Chassis, and the Extreme Networks logo are trademarks of Extreme Networks.

PACE is a trademark of 3Com Corporation. 3Com is a registered trademark of 3Com Corporation. All other brand and product names are registered trademarks or trademarks of their respective holders.

Contents

PREFACE

Introduction xv
Conventions xvi
Related Publications xvii

1 BLACKDIAMOND OVERVIEW

BlackDiamond 6800 Components 1-1
Summary of Features 1-3
- Full-Duplex 1-4
- Load Sharing 1-4
- Virtual LANs (VLANs) 1-4
- Spanning Tree Protocol (STP) 1-5
- Quality of Service (QoS) 1-5
- IP Unicast Routing 1-6
- IP Multicast Routing 1-6

Network Configuration Example 1-6
Factory Defaults 1-8

2 ACCESSING THE SYSTEM

Understanding the Command Syntax 2-2
- Syntax Helper 2-2
- Command Completion with Syntax Helper 2-2
- Abbreviated Syntax 2-3
- Command Shortcuts 2-3
- Numerical Ranges 2-3

Names 2-4
Symbols 2-4
Line-Editing Keys 2-5
Common Commands 2-6
Configuring Management Access 2-8
Default Accounts 2-9
Changing the Default Password 2-9
Creating a Management Account 2-10
Viewing Accounts 2-10
Deleting an Account 2-11
Methods of Managing the BlackDiamond 2-11
Using the Console Interface 2-12
Using Telnet 2-12
Connecting to Another Host Using Telnet 2-12
Configuring System IP Parameters 2-13
Using a BOOTP Server 2-13
Manually Configuring the IP Settings 2-14
Disconnecting a Telnet Session 2-15
Disabling Telnet Access 2-16
IP Host Configuration Commands 2-16
Using ExtremeWare Vista 2-17
Disabling Web Access 2-18
Using SNMP 2-18
Accessing System Agents 2-18
Supported MIBs 2-19
Configuring SNMP Settings 2-19
Displaying SNMP Settings 2-21
Resetting and Disabling SNMP 2-22
Checking Basic Connectivity 2-22
Ping 2-22
Traceroute 2-23

3 CONFIGURING SLOTS AND PORTS

Configuring a Slot 3-1
Specifying One or More Ports 3-3
Enabling and Disabling Ports 3-3
Configuring Port Speed and Duplex Setting 3-4
Turning Off Autonegotiation for a Gigabit Ethernet Port 3-4
Port Commands 3-4
Load Sharing 3-7
Configuring Load Sharing 3-7
Verifying the Load-Sharing Configuration 3-9
Port-Mirroring 3-9
Port-Mirroring Commands 3-10
Port-Mirroring Example 3-10

4 VIRTUAL LANS (VLANS)

Overview of Virtual LANs 4-1
Benefits 4-1
Types of VLANs 4-2
Port-Based VLANs 4-2
Spanning Systems with Port-Based VLANs 4-3
Tagged VLANs 4-6
Uses of Tagged VLANs 4-6
Assigning a VLAN Tag 4-6
Mixing Port-Based and Tagged VLANs 4-9
Generic VLAN Registration Protocol 4-9
GVRP Commands 4-11
Protocol-Based VLANs 4-11
Predefined Protocol Filters 4-12
Defining Protocol Filters 4-13
Deleting a Protocol Filter 4-14
Precedence of Tagged Packets Over Protocol Filters 4-14
VLAN Names 4-14
Default VLAN 4-15
Configuring VLANs on the BlackDiamond 4-15
VLAN Configuration Examples 4-17
Displaying VLAN Settings 4-18
Deleting VLANs 4-20

5 FORWARDING DATABASE (FDB)

Overview of the FDB 5-1
FDB Contents 5-1
FDB Entry Types 5-1
How FDB Entries Get Added 5-3
Associating a QoS Profile with an FDB Entry 5-3
Configuring FDB Entries 5-3
FDB Configuration Examples 5-4
Displaying FDB Entries 5-5
Removing FDB Entries 5-6

6 SPANNING TREE PROTOCOL (STP)

Overview of the Spanning Tree Protocol 6-1
Spanning Tree Domains 6-1
Defaults 6-2
STP Configurations 6-2
Configuring STP on the BlackDiamond 6-5
Configuration Example 6-8
Displaying STP Settings 6-8
Disabling and Resetting STP 6-9

7 QUALITY OF SERVICE (QOS)

Overview of Quality of Service 7-1
Building Blocks 7-1
QoS Mode 7-2
Default QoS Profiles 7-2
Traffic Groupings 7-3
Ingress Traffic Groupings 7-4
Egress Traffic Groupings 7-5
Precedence 7-5
Prioritization 7-6
Creating and Configuring a QoS Profile 7-6
Assigning a QoS Profile 7-7
Port Queue Monitor 7-7

Configuring QoS 7-8
Sample Ingress Mode QoS Configuration 7-9
Sample Egress Mode QoS Configuration 7-10
Displaying QoS Information 7-10
Resetting QoS 7-11

8 IP UNICAST ROUTING

Overview of IP Unicast Routing 8-1
Router Interfaces 8-2
Populating the Routing Table 8-3
Dynamic Routes 8-3
Static Routes 8-3
Multiple Routes 8-4
Proxy ARP 8-4
ARP-Incapable Devices 8-4
Proxy ARP Between Subnets 8-5
IP Multinetting 8-5
IP Multinetting Operation 8-6
IP Multinetting Examples 8-7
Configuring IP Unicast Routing 8-9
Verifying the IP Unicast Routing Configuration 8-9
Configuring DHCP/BOOTP Relay 8-10
Verifying the DHCP/BOOTP Relay Configuration 8-10
Using Extreme Standby Router Protocol 8-11
Configuring ESRP 8-11
IP Commands 8-12
Routing Configuration Example 8-16
Displaying Router Settings 8-18
Resetting and Disabling Router Settings 8-20

9 ROUTING PROTOCOLS

Overview 9-1
RIP Versus OSPF 9-2
Overview of RIP 9-3
Routing Table 9-3
Split Horizon 9-3
Poison Reverse 9-3

Triggered Updates 9-4
Route Advertisement of VLANs 9-4
RIP Version 1 versus RIP Version 2 9-4
Overview of OSPF 9-5
Link-State Database 9-5
Areas 9-5
Area 0 9-6
Stub Areas 9-6
Virtual Links 9-7
Unicast Route Distribution 9-9
Configuring RIP 9-9
RIP Configuration Example 9-11
Displaying RIP Settings 9-13
Resetting and Disabling RIP 9-14
Configuring OSPF 9-14
OSPF Configuration Example 9-17
Configuration for ABR1 9-19
Configuration for IR1 9-19
Displaying OSPF Settings 9-20
Resetting and Disabling OSPF Settings 9-20

10 IP MULTICAST ROUTING

Overview 10-1
DVMRP Overview 10-2
IGMP Overview 10-2
IGMP Snooping 10-2
Configuring IP Multicasting Routing 10-2
Configuration Example 10-5
Configuration for IR1 10-7
Displaying IP Multicast Routing Settings 10-7
Deleting and Resetting IP Multicast Settings 10-8

11 STATUS MONITORING AND STATISTICS

Status Monitoring 11-1
Slot Diagnostics 11-8
Port Statistics 11-9
Port Errors 11-10
Port Monitoring Display Keys 11-11
Logging 11-11
- Local Logging 11-13
  - Real-Time Display 11-13
- Remote Logging 11-14
- Logging Commands 11-14

RMON 11-16
- About RMON 11-17
- RMON Features of the BlackDiamond 11-17
  - Statistics 11-17
  - History 11-17
  - Alarms 11-18
  - Events 11-18
- RMON and the BlackDiamond 11-18
- Event Actions 11-19

12 USING EXTREMEWARE VISTA

Enabling and Disabling Web Access 12-1
Setting Up Your Browser 12-2
Accessing ExtremeWare Vista 12-3
Navigating ExtremeWare Vista 12-3
- Task Frame 12-4
- Content Frame 12-4
  - Browser Controls 12-4
  - Status Messages 12-5
- Standalone Buttons 12-5

Saving Changes 12-5
Filtering Information 12-6
Do a GET When Configuring a VLAN 12-6
Sending Screen Output to Extreme Networks 12-7

13 SOFTWARE UPGRADE AND BOOT OPTIONS
Downloading a New Image 13-1
Rebooting the System 13-2
Saving Configuration Changes 13-3
Returning to Factory Defaults 13-3
Using TFTP to Upload the Configuration 13-4
Using TFTP to Download the Configuration 13-5
Boot Option Commands 13-5

A SUPPORTED STANDARDS

B TROUBLESHOOTING
LEDs 9
Using the Command-Line Interface 10
VLANs 12
STP 14

INDEX

Figures

1-1 BlackDiamond 6800 system 1-2
1-2 BlackDiamond network configuration example 1-7
4-1 Example of a port-based VLAN 4-3
4-2 Single port-based VLAN spanning two switches 4-4
4-3 Two port-based VLANs spanning two BlackDiamond systems 4-5
4-4 Physical diagram of tagged and untagged traffic 4-7
4-5 Logical diagram of tagged and untagged traffic 4-8
4-6 Network example using GVRP 4-9
4-7 Protocol-based VLANs 4-12
6-1 Multiple Spanning Tree Domains 6-3
6-2 Tag-based STP configuration 6-4
8-1 Routing between VLANs 8-2
8-2 Unicast routing configuration example 8-17
9-1 Stub area 9-7
9-2 Virtual link for stub area 9-8
9-3 Virtual link providing redundancy 9-8
9-4 RIP configuration example 9-12
9-5 OSPF configuration example 9-17
10-1 IP multicast routing configuration example 10-6

Tables

1 Notice Icons xvi
2 Text Conventions xvi
1-1 BlackDiamond Factory Defaults 1-8
2-1 Command Syntax Symbols 2-4
2-2 Line-Editing Keys 2-5
2-3 Common Commands 2-6
2-4 Default Accounts 2-9
2-5 IP Host Configuration Commands 2-16
2-6 Supported MIBs 2-19
2-7 SNMP Configuration Commands 2-20
2-8 SNMP Reset and Disable Commands 2-22
2-9 Ping Command Parameters 2-22
3-1 Port Commands 3-5
3-2 Port Combinations for the G4X Module 3-8
3-3 Port Combinations for the G6X Module 3-8
3-4 Port Combinations for the F32T and F32F Modules 3-8
3-5 Port-Mirroring Configuration Commands 3-10
4-1 GVRP Commands 4-11
4-2 VLAN Configuration Commands 4-15
4-3 VLAN Delete and Reset Commands 4-20
5-1 FDB Configuration Commands 5-3
5-2 Removing FDB Entry Commands 5-6
6-1 STP Configuration Commands 6-6
6-2 STP Disable and Reset Commands 6-9
7-1 Default QoS Profiles 7-3
7-2 802.1p Values and Associated QoS Profiles 7-4

7-3 PQM Commands 7-8
7-4 QoS Configuration Commands 7-8
8-1 ESRP Commands 8-11
8-2 Basic IP Commands 8-12
8-3 Route Table Configuration Commands 8-14
8-4 ICMP Configuration Commands 8-15
8-5 Router Show Commands 8-18
8-6 Router Reset and Disable Commands 8-20
9-1 RIP Configuration Commands 9-9
9-2 RIP Show Commands 9-13
9-3 RIP Reset and Disable Commands 9-14
9-4 OSPF Configuration Commands 9-14
9-5 OSPF Show Commands 9-20
9-6 OSPF Reset and Disable Commands 9-20
10-1 IP Multicast Routing Configuration Commands 10-3
10-2 IGMP Configuration Commands 10-4
10-3 IP Multicast Routing Show Commands 10-7
10-4 IP Multicast Routing Reset and Disable Commands 10-8
11-1 Switch Monitoring Commands 11-1
11-2 Port Monitoring Display Keys 11-11
11-3 Fault Levels Assigned by the System 11-12
11-4 Fault Log Subsystems 11-12
11-5 Logging Commands 11-15
11-6 Event Actions 11-19
12-1 Multi-Select List Box Key Definitions 12-4
13-1 Boot Option Commands 13-5

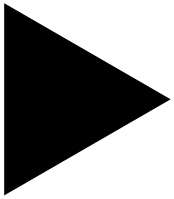

Preface

This Preface provides an overview of this guide, describes guide conventions, and lists other publications that may be useful.

INTRODUCTION

This guide provides the required information to configure the BlackDiamond™ system and software.

This guide is intended for use by network administrators who are responsible for installing and setting up network equipment. It assumes a basic working knowledge of the following:

- Local Area Networks (LANs)
- Ethernet concepts
- Ethernet switching and bridging concepts
- Routing concepts
- Internet Protocol (IP) concepts
- Routing Information Protocol (RIP) and Open Shortest Path First (OSPF)
- IP Multicast concepts
- Distance Vector Multicast Routing Protocol (DVMRP) concepts
- Simple Network Management Protocol (SNMP)

If the information in the "Release Notes" shipped with your switch differs from the information in this guide, follow the "Release Notes."

CONVENTIONS

Table 1 and Table 2 list conventions that are used throughout this guide.

Table 1: Notice Icons

Icon	Notice Type	Alerts you to...
	Note	Important features or instructions.
	Caution	Risk of personal injury, system damage, or loss of data.
	Warning	Risk of severe personal injury.

Table 2: Text Conventions

Convention	Description
`Screen displays`	This typeface indicates command syntax, or represents information as it appears on the screen.
`Screen displays bold`	This typeface indicates how you would type a particular command.
The words "enter" and "type"	When you see the word "enter" in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says "type."
[Key] names	Key names are written with brackets, such as [Return] or [Esc]. If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press [Ctrl]+[Alt]+[Del].
Words in *italicized* type	Italics emphasize a point or denote new terms at the place where they are defined in the text.

The command syntax is explained in Chapter 2.

RELATED PUBLICATIONS

The BlackDiamond documentation set includes the following:

- *BlackDiamond Hardware Installation Guide*
- *BlackDiamond Quick Reference Guide*
- BlackDiamond “Release Notes”

Documentation for Extreme Networks products is available on the World Wide Web at the following location:

- Extreme Networks home page http://www.extremenetworks.com/

1 BlackDiamond Overview

The BlackDiamond 6800 is a chassis-based system that is designed to be placed in the core of your network. The BlackDiamond system is flexible and scalable, making it easy for you to meet the changing requirements of your network. The combination of BlackDiamond systems and Summit switches delivers a consistent end-to-end network solution that provides non-blocking architecture, wire-speed switching, wire-speed IP routing, and policy-based Quality of Service.W

This chapter describes the following:

- BlackDiamond 6800 components
- BlackDiamond 6800 features
- How to use the BlackDiamond 6800 system in your network configuration
- Factory default settings

BLACKDIAMOND 6800 COMPONENTS

The BlackDiamond 6800 system, shown in Figure 1-1, consists of the following components:

- One 10-slot chassis with backplane
- Eight I/O module slots, labeled Slot 1 through Slot 8
- Two MSM module slots, labeled Slot A and Slot B
- As many as two load-sharing, fully redundant power supplies
- One hot-swappable fan tray (accessed from the rear of the unit)

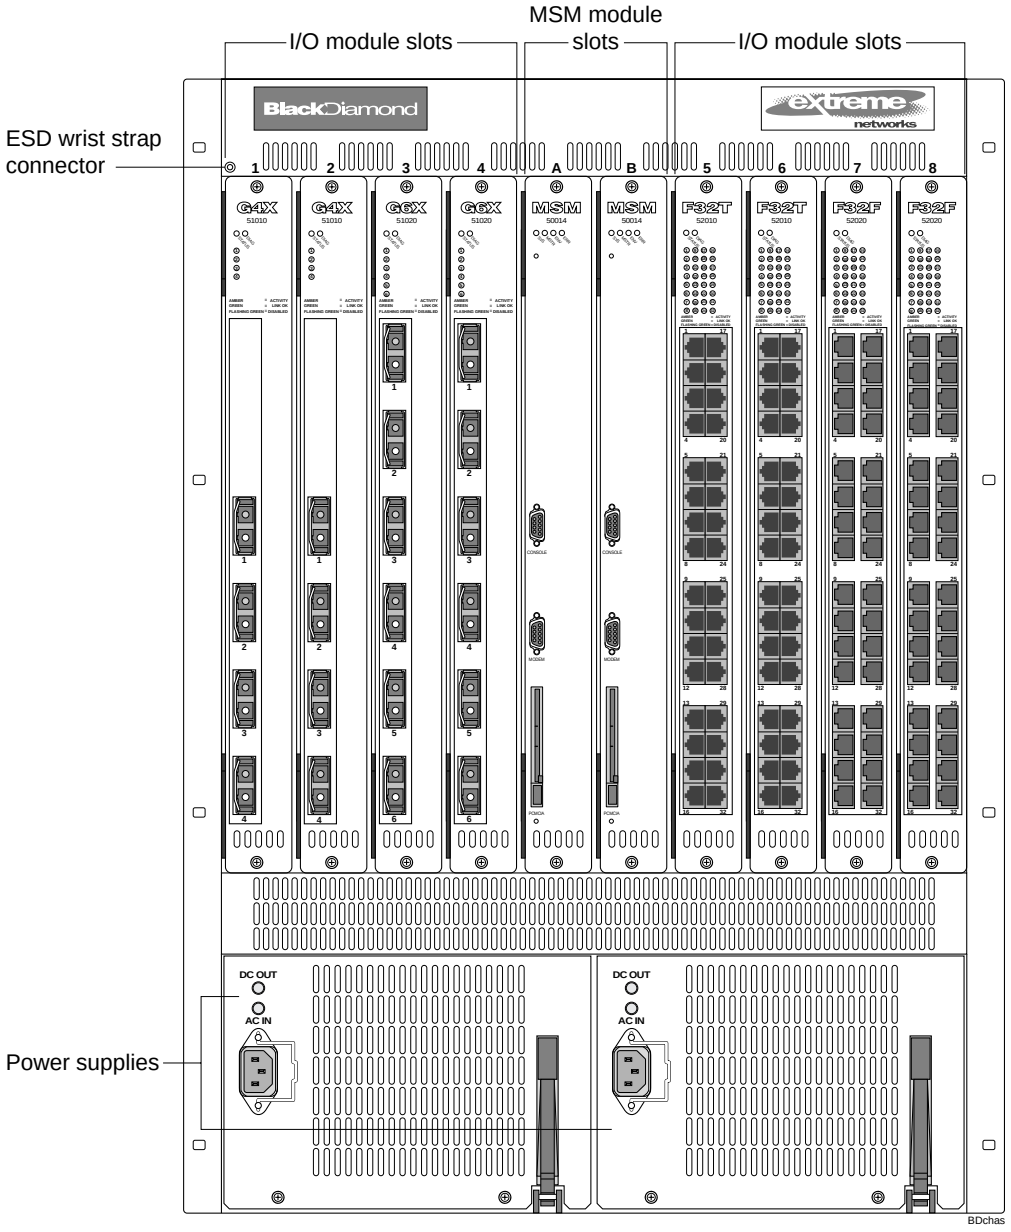

Figure 1-1: BlackDiamond 6800 system

SUMMARY OF FEATURES

The features of the BlackDiamond 6800 include the following:

- A 10-slot chassis that can be populated with up to 8 input/output (I/O) modules and 2 Management Switch Fabric Modules (MSMs)
- MSM and I/O modules are hot-swappable
- I/O modules include Gigabit Ethernet or 10/100 Mbps Ethernet ports
- Redundant, load-sharing, hot-swappable power supplies
- Field-replaceable fan tray
- Up to 256 switched 10/100 Mbps Ethernet ports
- Up to 48 switched Gigabit Ethernet ports
- Fully nonblocking operation
  - All ports transmit and receive packets at wire speed
- Autonegotiation for half- or full-duplex operation (10BASE-T/100BASE-TX ports, only)
- Load sharing on multiple ports
- Virtual local area networks (VLANs) including support for IEEE 802.1Q and IEEE 802.1p
- Spanning Tree Protocol (STP) (IEEE 802.1D) with multiple STP domains
- Policy-Based Quality of Service (PB-QoS)
- Wire-speed Internet Protocol (IP) routing
- IP Multinetting
- DHCP/BOOTP Relay
- Routing Information Protocol (RIP) version 1 and RIP version 2
- Open Shortest Path First (OSPF) routing protocol
- Wire-speed IP multicast routing support
- IGMP snooping to control IP multicast traffic
- Distance Vector Multicast Routing Protocol (DVMRP)
- Console command-line interface (CLI) connection
- Telnet CLI connection
- ExtremeWare Vista Web-based management interface

- Simple Network Management Protocol (SNMP) support
- Remote Monitoring (RMON)
- Traffic mirroring for all ports

For more information on BlackDiamond components, refer to the BlackDiamond Hardware Installation Guide.

FULL-DUPLEX

The BlackDiamond provides full-duplex support for all ports. Full-duplex allows frames to be transmitted and received simultaneously and, in effect, doubles the bandwidth available on a link. All 10/100 Mbps ports on the BlackDiamond autonegotiate for half- or full-duplex operation.

LOAD SHARING

Load sharing with the BlackDiamond allows you to increase bandwidth and resilience by using a group of ports to carry traffic in parallel between systems. The sharing algorithm allows the BlackDiamond system to use multiple ports as a single logical port. For example, VLANs see the load-sharing group as a single virtual port. The algorithm also guarantees packet sequencing between clients.

For information on load sharing, refer to Chapter 3.

VIRTUAL LANS (VLANS)

The BlackDiamond has a VLAN feature that enables you to construct your broadcast domains without being restricted by physical connections. Up to 255 VLANs can be defined on the BlackDiamond. A VLAN is a group of location- and topology-independent devices that communicate as if they were on the same physical local area network (LAN).

Implementing VLANs on your network has the following three advantages:

- It helps to control broadcast traffic. If a device in VLAN *marketing* transmits a broadcast frame, only VLAN *marketing* devices receive the frame.
- It provides extra security. Devices in VLAN *marketing* can only communicate with devices on VLAN *sales* using routing services.

- It eases the change and movement of devices on networks. If a device in VLAN *marketing* is moved to a port in another part of the network, all you must do is specify that the new port belongs to VLAN *marketing*.

For more information on VLANs, refer to Chapter 4.

SPANNING TREE PROTOCOL (STP)

The BlackDiamond supports the IEEE 802.1D Spanning Tree Protocol (STP), which is a bridge-based mechanism for providing fault tolerance on networks. STP enables you to implement parallel paths for network traffic, and ensure the following:

- Redundant paths are disabled when the main paths are operational.
- Redundant paths are enabled if the main traffic paths fail.

The BlackDiamond supports up to 64 Spanning Tree Domains (STPDs).

For more information on STP, refer to Chapter 6.

QUALITY OF SERVICE (QOS)

The BlackDiamond has Policy-Based Quality of Service (QoS) features that enable you to specify service levels for different traffic groups. By default, all traffic is assigned the "normal" QoS policy profile. If needed, you can create other QoS policies and apply them to different traffic types so that they have different guaranteed minimum bandwidth, maximum bandwidth, and priority.

For more information on Quality of Service, refer to Chapter 7.

IP UNICAST ROUTING

The BlackDiamond can route IP traffic between the VLANs that are configured as virtual router interfaces. Both dynamic and static IP routes are maintained in the routing table. The following routing protocols are supported:

- RIP version 1
- RIP version 2
- OSPF

For more information on IP unicast routing, refer to Chapter 8.

IP MULTICAST ROUTING

The BlackDiamond can use IP multicasting to allow a single IP host to transmit a packet to a group of IP hosts. The BlackDiamond supports multicast routes that are learned by way of the Distance Vector Multicast Routing Protocol (DVMRP).

For more information on IP multicast routing, refer to Chapter 10.

NETWORK CONFIGURATION EXAMPLE

The BlackDiamond can be used as a core network switch, a segment switch, or a server switch. When used as a core network switch, the BlackDiamond performs collapsed-backbone routing and switching at wire-speed, allowing it to handle the growing influx of intranetwork and Internet traffic that travels between multiple networks.

When used as a high-density segment switch in the wiring closet, the BlackDiamond can aggregate the hierarchy of legacy hubs and switches that support desktop connections. When the BlackDiamond is used to perform segment switching at Layer 3, local traffic stays within the subnet and does not needlessly enter the core.

When used as a server switch, the BlackDiamond supports 10/100 Mbps segments and Gigabit Ethernet segments, handling traffic to and from servers in the data center. This application ensures sufficient bandwidth between servers and to segments. Conversely, the BlackDiamond can support Gigabit Ethernet segments and 10/100 Mbps server links to optimize client-to-server traffic. Multiple 10 Mbps, 100 Mbps, and 1,000 Mbps switch ports can be trunked into one logical link between the data centers, the network cores, and high-performance servers.

Figure 1-2 shows an example of the BlackDiamond used in a network configuration.

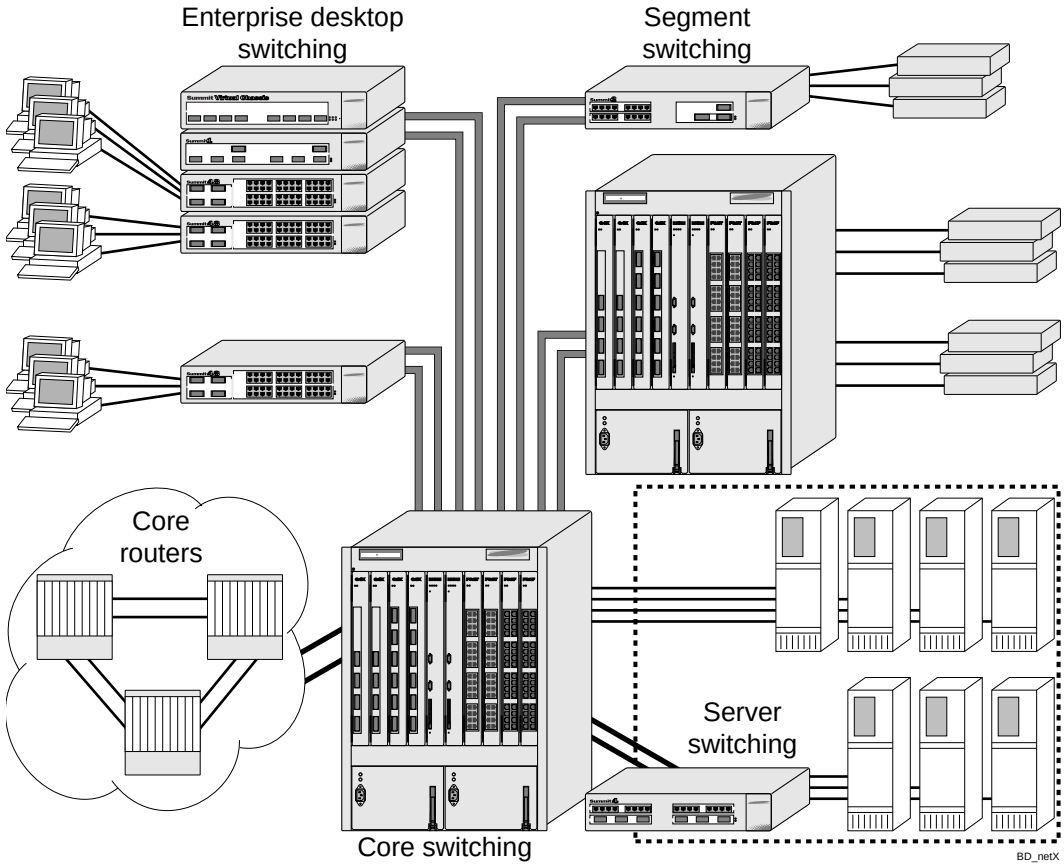

Figure 1-2: BlackDiamond network configuration example

FACTORY DEFAULTS

Table 1-1 shows factory defaults for the BlackDiamond features.

Table 1-1: BlackDiamond Factory Defaults

Item	Default Setting
Port status	Enabled on all ports
Serial or Telnet user account	*admin* with no password and *user* with no password
Console port configuration	9600 baud, 8 data bits, 1 stop bit, no parity, XON/XOFF flow control enabled
Web network management	Enabled
SNMP read community string	*public*
SNMP write community string	*private*
RMON history session	Enabled
RMON alarms	Disabled
BOOTP	Enabled on the default VLAN (*default*)
QoS	All traffic is part of the default queue in ingress mode
QoS monitoring	Automatic roving
802.1p priority	Recognition enabled
802.3x flow control	Enabled on Gigabit Ethernet ports
Virtual LANs	One VLAN named *default*; all ports belong to the default VLAN; the default VLAN belongs to the STPD named *s0*
802.1Q tagging	All packets are untagged on the default VLAN (*default)*
Spanning Tree Protocol	Disabled for the switch; enabled for each port in the STPD
Forwarding database aging period	300 seconds (5 minutes)
IP Routing	Disabled
RIP	Disabled
OSPF	Disabled
IP multicast routing	Disabled
IGMP snooping	Enabled
DVMRP	Disabled
GVRP	Disabled
Port mirroring	Disabled

2 Accessing The System

This chapter provides the following required information to begin managing the BlackDiamond:

- Understanding the command syntax
- Line-editing commands
- Command history substitution
- Configuring the system for management
- System management methods
- Configuring SNMP
- Checking basic connectivity

In order for configuration changes to be retained through a power cycle or reboot, you must issue a SAVE command after you have made the change. For more information on the SAVE command, refer to Chapter 13.

UNDERSTANDING THE COMMAND SYNTAX

This section describes the steps to take when entering a command. Refer to the sections that follow for detailed information on using the command-line interface.

To use the command-line interface, follow these steps:

1 When entering a command at the prompt, ensure that you have the appropriate privilege level.

Most configuration commands require you to have the administrator privilege level.

2 Enter the command name.

If the command does not include a parameter or values, skip to Step 3. If the command requires more information, continue to Step 2a.

a If the command includes a parameter, enter the parameter name and values.

The value part of the command specifies how you want the parameter to be set. Values include numerics, strings, or addresses, depending on the parameter.

3 After entering the complete command, press [Return].

If an asterisk (*) appears in front of the command-line prompt, it indicates that you have outstanding configuration changes that have not been saved. For more information on saving configuration changes, refer to Chapter 13.

SYNTAX HELPER

The command-line interface has a built-in syntax helper. If you are unsure of the complete syntax for a particular command, enter as much of the command as possible and press [Return]. The syntax helper provides a list of options for the remainder of the command.

The syntax helper also provides assistance if you have entered an incorrect command.

COMMAND COMPLETION WITH SYNTAX HELPER

The BlackDiamond provides command completion by way of the [Tab] key. If you enter a partial command, pressing the [Tab] key posts a list of available options, and places the cursor at the end of the command.

ABBREVIATED SYNTAX

Abbreviated syntax is the shortest, most unambiguous, allowable abbreviation of a command or parameter. Typically, this is the first three letters of the command.

COMMAND SHORTCUTS

All named components of the system configuration must have a unique name. Components are named using the `create` command. When you enter a command to configure a named component, you do not need to use the keyword of the component. For example, to create a VLAN, you must enter a unique VLAN name:

```
create vlan engineering
```

Once you have created the VLAN with a unique name, you can then eliminate the keyword `vlan` from all other commands that require the name to be entered. For example, instead of entering the command:

```
config vlan engineering delete port 1:3,4:6
```

you could enter the following shortcut:

```
config engineering delete port 1:3,4:6
```

NUMERICAL RANGES

Commands that require you to enter one or more port numbers use the parameter `<slot:portlist>` in the syntax. A `<portlist>` can be one port on a particular slot. For example,

```
port 3:1
```

A `<portlist>` can be a range of numbers. For example,

```
port 3:1-3:3
```

You can add additional slot and port numbers to the list, separated by a comma:

```
port 3:1,4:8,6:10
```

You can specify all ports on a particular slot. For example,

```
port 3:*
```

indicates all ports on slot 3.

You can specify a range of slots and ports. For example,

```
port 2:3-4:5
```

indicates slot 2, port 3 through slot 4, port 5.

NAMES

All named components of the system configuration must have a unique name. Names must begin with an alphabetical character and are delimited by whitespace, unless enclosed in quotation marks.

SYMBOLS

You may see a variety of symbols shown as part of the command syntax. These symbols explain how to enter the command, and you do not type them as part of the command itself. Table 2-1 summarizes command syntax symbols.

Table 2-1: Command Syntax Symbols

Symbol	Description
angle brackets < >	Enclose a variable or value. You must specify the variable or value. For example, in the syntax `config vlan <name> ipaddress <ip_address>` you must supply a VLAN name for `<name>` and an address for `<ip_address>` when entering the command. Do not type the angle brackets.
square brackets [ ]	Enclose a required value or list of required arguments. One or more values or arguments can be specified. For example, in the syntax `disable vlan [<name> | all]` you must specify either the VLAN name for `<name>`, or the keyword `all` when entering the command. Do not type the square brackets.
vertical bar |	Separates mutually exclusive items in a list, one of which must be entered. For example, in the syntax `config snmp community [readonly | readwrite] <string>` you must specify either the read or write community string in the command. Do not type the vertical bar.

Table 2-1: Command Syntax Symbols (continued)

Symbol	Description
braces { }	Enclose an optional value or a list of optional arguments. One or more values or arguments can be specified. For example, in the syntax `show vlan {<name> | all}` you can specify either a particular VLAN or the keyword `all`. If you do not specify an argument, the command will show all VLANs. Do not type the braces.

LINE-EDITING KEYS

Table 2-2 describes the line-editing keys available using the command-line interface.

Table 2-2: Line-Editing Keys

Key(s)	Description
Backspace	Deletes character to the left of cursor and shifts remainder of line to left.
Delete or [Ctrl] + D	Deletes character under cursor and shifts remainder of line to left.
[Ctrl] + K	Deletes characters from under cursor to the end of the line.
Insert	Toggles on and off. When toggled on, inserts text and shifts previous text to right.
Left Arrow	Moves cursor to left.
Right Arrow	Moves cursor to right.
Home or [Ctrl] + A	Moves cursor to first character in line.
End or [Ctrl] + E	Moves cursor to last character in line.
[Ctrl] + L	Clears the screen and movers the cursor to the beginning of the line.
[Ctrl] + U	Clears all characters typed from the cursor to the beginning of the line.
[Ctrl] + W	Deletes the previous word.
Up Arrow	Displays the previous command in the command history buffer and places cursor at end of command.
Down Arrow	Displays the next command in the command history buffer and places cursor at end of command.

COMMAND HISTORY

The BlackDiamond "remembers" the last 49 commands you enter. You can display a list of these commands by using the following command:

```
history
```

COMMON COMMANDS

Table 2-3 describes common commands used to manage the system. Commands specific to a particular feature are described in the other chapters of this guide.

Table 2-3: Common Commands

Command	Description
create account [admin | user] <username> {<password>}	Creates a user account.
create vlan <name>	Creates a VLAN.
config account <username> {<password>}	Configures a user account password. Passwords must have a minimum of 4 characters and can have a maximum of 12 characters. User names and passwords are case-sensitive.
config banner	Configures the banner string. You can enter up to 24 rows of 80-column text that is displayed before the login prompt of each session. Press [Return] at the beginning of a line to terminate the command and apply the banner. To clear the banner, press [Return] at the beginning of the first line.
config slot <slot> module [f32t | f32f | g4x | g6x]	Configures a slot for a particular I/O module card.
config port [<slot:portlist> | all] auto off {speed [10 | 100]} duplex [half | full]	Manually configures the port speed and duplex setting of one or more ports.
config time <date> <time>	Configures the system date and time. The format is as follows: `mm/dd/yyyy hh:mm:ss` The time uses a 24-hour clock format. You cannot set the year past 2023.

Table 2-3: Common Commands (continued)

Command	Description
config vlan <name> ipaddress <ip_address> {<mask>}	Configures an IP address and subnet mask for a VLAN.
enable bootp vlan [<name> | all]	Enables BOOTP for one or more VLANs.
enable idletimeout	Enables a timer that disconnects all sessions (both Telnet and console) after 20 minutes of inactivity. The default setting is disabled.
clear session <number>	Terminates a Telnet session from the system.
disable bootp vlan [<name> | all]	Disables BOOTP for one or more VLANs.
disable idletimeout	Disables the timer that disconnects all sessions. Once disabled, console sessions remain open until the system is rebooted. Telnet sessions remain open until you close the Telnet client.
disable port [<slot:portlist> | all]	Disables a port.
disable telnet	Disables Telnet access to the system.
disable web	Disables Web access to the system.
delete account <username>	Deletes a user account.
delete vlan <name>	Deletes a VLAN.
unconfig switch {all}	Resets all system parameters (with the exception of defined user accounts) to the factory defaults. If you specify the keyword `all`, the user account information is reset as well.
show banner	Displays the user-configured banner.

CONFIGURING MANAGEMENT ACCESS

The BlackDiamond supports the following two level levels of management:

- User
- Administrator

A user-level account has viewing access to all manageable parameters, with the exception of the following:

- User account database
- SNMP community strings

A user-level account can use the `ping` command to test device reachability, and change the password assigned to the account name. If you have logged on with user capabilities, the command-line prompt ends with a (>) sign. For example:

```
BD6800:2>
```

An administrator-level account can view and change all system parameters. It can also add and delete users, and change the password associated with any account name. The administrator can disconnect a management session that has been established by way of a Telnet connection. If this happens, the user logged on by way of the Telnet connection is notified that the session has been terminated.

If you have logged on with administrator capabilities, the command-line prompt ends with a (#) sign. For example:

```
BD6800:18#
```

The prompt text is taken from the SNMP `sysname` setting. The number that follows the colon indicates the sequential line/command number.

If an asterisk (*) appears in front of the command-line prompt, it indicates that you have outstanding configuration changes that have not been saved. For example:

```
*BD6800:19#
```

For more information on saving configuration changes, refer to Chapter 13.

DEFAULT ACCOUNTS

By default, the system is configured with two accounts, as shown in Table 2-4.

Table 2-4: Default Accounts

Account Name	Access Level
admin	This user can access and change all manageable parameters. The admin account cannot be deleted.
user	This user can view (but not change) all manageable parameters, with the following exceptions: ■ This user cannot view the user account database. ■ This user cannot view the SNMP community strings. This user has access to the `ping` command.

CHANGING THE DEFAULT PASSWORD

Default accounts do not have passwords assigned to them. Passwords must have a minimum of 4 characters and can have a maximum of 12 characters.

User names and passwords are case-sensitive.

To add a password to the default admin account, follow these steps:

1. Log in to the system using the name *admin*.
2. At the password prompt, press [Return].
3. Add a default admin password by typing the following:

   `config account admin`

4. Enter the new password at the prompt.
5. Re-enter the new password at the prompt.

To add a password to the default user account, follow these steps:

1. Log in to the system using the name *admin*.
2. At the password prompt, press [Return], or enter the password that you have configured for the *admin* account.
3. Add a default user password by typing the following:

   `config account user`

4 Enter the new password at the prompt.

5 Re-enter the new password at the prompt.

If you forget your password while logged out of the command-line interface, contact your local technical support representative, who will advise on your next course of action.

CREATING A MANAGEMENT ACCOUNT

The system can have a total of 16 management accounts. You can use the default names (*admin* and *user*), or you can create new names and passwords for the accounts. Passwords must have a minimum of 4 characters and can have a maximum of 12 characters.

The account name admin *cannot be deleted.*

To create a new account, follow these steps:

1 Log in to the system as *admin*.

2 At the password prompt, press [Return], or enter the password that you have configured for the *admin* account.

3 Add a new user by using the following command:

```
create account [admin | user] <username>
```

4 Enter the password at the prompt.

5 Re-enter the password at the prompt.

VIEWING ACCOUNTS

To view the accounts that have been created, you must have administrator privileges. Use the following command to see the accounts:

```
show account
```

Output from the `show accounts` command is as follows:

```
#show account

User Name      Access  LoginOK  Failed  Session
-------------  ------  -------  ------  --------
       admin    R/W       0         0
        user    RO        0         0
```

DELETING AN ACCOUNT

To delete a account, you must have administrator privileges. Use the following command to delete an account:

```
delete account <username>
```

METHODS OF MANAGING THE BLACKDIAMOND

You can manage the BlackDiamond using the following methods:

- Access the command-line interface by connecting a terminal (or workstation with terminal-emulation software) to the MSM port labelled *console*.
- Access the command-line interface over a TCP/IP network using a Telnet connection.
- Access the Web interface over a TCP/IP network, using a standard Web browser (such as Netscape Navigator 3.0 or greater, or Microsoft Internet Explorer 3.0 or greater).
- Use an SNMP Network Manager over a network running the IP protocol.

The system can support up to multiple user sessions concurrently, as follows:

- Two console sessions
- Eight Telnet sessions
- One Web session

USING THE CONSOLE INTERFACE

The command-line interface built into the system is accessible by way of the 9-pin, RS-232 port labelled *console*, located on either of the MSMs.

For more information on the console port pinouts, refer to the BlackDiamond Hardware Installation Guide.

Once the connection is established, you will see the system prompt and you may log in.

USING TELNET

Any workstation with a Telnet facility should be able to communicate with the system over a TCP/IP network.

Up to eight active Telnet sessions can access the system concurrently. If `idle timeouts` are enabled, the Telnet connection will time out after 20 minutes of inactivity. If a connection to a Telnet session is lost inadvertently, the system terminates the session within two hours.

Before you can start a Telnet session, you must set up the IP parameters described in the section "Configuring System IP Parameters," later in this chapter. Telnet is enabled by default.

To open the Telnet session, you must specify the IP address of the device that you want to manage. Check the user manual supplied with the Telnet facility if you are unsure of how to do this.

Once the connection is established, you will see the system prompt and you may log in.

CONNECTING TO ANOTHER HOST USING TELNET

You can Telnet from the current command-line interface session to another host using the following command:

```
telnet <ipaddress> {<port_number>}
```

If the TCP port number is not specified, the Telnet session defaults to port 23. Only VT100 emulation is supported.

CONFIGURING SYSTEM IP PARAMETERS

To manage the system by way of a Telnet connection or by using an SNMP Network Manager, you must first configure the system IP parameters.

USING A BOOTP SERVER

If you are using IP and you have a BOOTP server set up correctly on your network, you must add the following information to the BOOTP server:

- System Media Access Control (MAC) address
- IP address
- Subnet address mask (optional)

The system MAC address is found on the rear label of the system.

Once this is done, the IP address and subnetwork mask for the system will be downloaded automatically. You can then start managing the system without further configuration.

You can enable BOOTP on a per-VLAN basis by using the following command:

```
enable bootp vlan [<name> | all]
```

By default, BOOTP is enabled on the *default* VLAN.

If you configure the system to use BOOTP, the system IP address is not retained through a power cycle, even if the configuration has been saved. To retain the IP address through a power cycle, you must configure the IP address of the VLAN using the command-line interface, Telnet, or Web interface.

All VLANs within a BlackDiamond system that are configured to use BOOTP to get their IP address use the same MAC address. Therefore, if you are using BOOTP relay through a router, the BOOTP server must be capable of differentiating its relay based on the gateway portion of the BOOTP packet.

For more information on DHCP/BOOTP relay, refer to Chapter 8.

MANUALLY CONFIGURING THE IP SETTINGS

If you are using IP without a BOOTP server, you must enter the IP parameters for the system in order for the SNMP Network Manager, Telnet software, or Web interface to communicate with the device. To assign IP parameters to the system, you must do the following:

- Log in to the system with administrator privileges.
- Assign an IP address and subnetwork mask to a VLAN.

  The BlackDiamond comes configured with a default VLAN named *default*. To use Telnet or an SNMP Network Manager, you must have at least one VLAN on the system, and it must be assigned an IP address and subnetwork mask. IP addresses are always assigned to a VLAN. The BlackDiamond can be assigned multiple IP addresses.

For information on creating and configuring VLANs, refer to Chapter 4.

To manually configure the IP settings, perform the following steps:

1. Connect a terminal or workstation running terminal-emulation software to the console port.
2. At your terminal, press [Return] one or more times until you see the login prompt.
3. At the login prompt, enter your user name and password. Note that they are both case-sensitive. Ensure that you have entered a user name and password with administrator privileges.

   — If you are logging in for the first time, use the default user name *admin* to log in with administrator privileges. For example:

   ```
   login: admin
   ```

   Administrator capabilities enable you to access all system functions. The default user names have no passwords assigned.

   — If you have been assigned a user name and password with administrator privileges, enter them at the login prompt.
4. At the password prompt, enter the password and press [Return].

   When you have successfully logged in to the system, the command-line prompt displays the name of the system in its prompt.

5 Assign an IP address and subnetwork mask for the default VLAN by using the following command:

```
config vlan <name> ipaddress <ipaddress> {<subnet_mask>}
```

For example:

```
config vlan default ipaddress 123.45.67.8 255.255.255.0
```

Your changes take effect immediately.

6 Configure the default route for the system using the following command:

```
config iproute add default <ipaddress> {<metric>}
```

For example:

```
config iproute add default 123.45.67.1
```

7 Save your configuration changes so that they will be in effect after the next system reboot, by typing

```
save
```

For more information on saving configuration changes, refer to Chapter 13.

8 When you are finished using the facility, log out of the system by typing

```
logout
```

DISCONNECTING A TELNET SESSION

An administrator-level account can disconnect a management session that has been established by way of a Telnet connection. If this happens, the user logged in by way of the Telnet connection is notified that the session has been terminated.

To terminate a Telnet session, follow these steps:

1 Log in to the system with administrator privileges.

2 Determine the session number of the session you want to terminate by typing

```
show session
```

Sample output from the `show session` command is as follows:

```
show session:

0  Wed Sep 17 20:48:38 1997   admin  console serial
4  Wed Sep 17 21:52:16 1997   admin  telnet 192.208.37.26
```

The session number is the first number on the line of the `show session` output.

3 Terminate the session by using the following command:

```
clear session <session_number>
```

DISABLING TELNET ACCESS

By default, Telnet services are enabled on the system. You can choose to disable Telnet by entering

`disable telnet`

To re-enable Telnet on the system, at the console port enter

`enable telnet`

You must be logged in as an administrator to enable or disable Telnet.

IP HOST CONFIGURATION COMMANDS

Table 2-5 describes the commands that are used to configure IP settings on the system.

Table 2-5: IP Host Configuration Commands

Command	Description
config iparp add <ipaddress> <mac_address>	Adds a permanent entry to the Address Resolution Protocol (ARP) table. Specify the IP address and MAC address of the entry.
config iparp delete <ipaddress>	Deletes an entry from the ARP table. Specify the IP address of the entry.
clear iparp [<ipaddress> | vlan <name> | all]	Removes dynamic entries in the IP ARP table. Permanent IP ARP entries are not affected.

Table 2-5: IP Host Configuration Commands (continued)

Command	Description
config iproute add <ipaddress> <mask> <gateway> {<metric>}	Adds a static address to the routing table. Use a value of 255.255.255.255 for `mask` to indicate a host entry.
config iproute delete <ipaddress> <mask> <gateway>	Deletes a static address from the routing table.
config iproute add default <gateway> {<metric>}	Adds a default gateway to the routing table. A default gateway must be located on a configured IP interface. If no metric is specified, the default metric of 1 is used.
config iproute delete default <gateway>	Deletes a default gateway from the routing table.
show ipconfig {vlan <name> | all}	Displays configuration information for one or more VLANs.
show ipstats {vlan [<name> | all]}	Displays IP statistics for the CPU of the system.
show iparp {<ipaddress | vlan <name> | all | permanent}	Displays the IP ARP table. You can filter the display by IP address, VLAN, or permanent entries.

USING EXTREMEWARE VISTA

ExtremeWare Vista is device-management software running in the BlackDiamond that enables you to access the system over a TCP/IP network using a standard Web browser. Any properly configured standard Web browser that supports frames (such as Netscape Navigator 3.0 or Microsoft Internet Explorer 3.0) can manage the system over a TCP/IP network.

For more information on assigning an IP address, refer to "Configuring System IP Parameters," on page 2-13.

The default home page of the BlackDiamond system can be accessed using the following command:

```
http://<ipaddress>
```

When you access the home page of the system, you are presented with the Logon screen.

For more information on using ExtremeWare Vista, refer to Chapter 12.

DISABLING WEB ACCESS

By default, Web access is enabled on the BlackDiamond. To disable it, enter the following command:

```
disable web
```

To re-enable Web access, enter the following command:

```
enable web
```

Reboot the system in order for these changes to take effect.

For more information on rebooting the system, refer to Chapter 13.

USING SNMP

Any Network Manager running the Simple Network Management Protocol (SNMP) can manage the system, provided the Management Information Base (MIB) is installed correctly on the management station. Each Network Manager provides its own user interface to the management facilities.

The following sections describe how to get started if you want to use an SNMP manager. It assumes you are already familiar with SNMP management. If not, refer to the following publication:

The Simple Book
by Marshall T. Rose
ISBN 0-13-8121611-9
Published by Prentice Hall

ACCESSING SYSTEM AGENTS

In order to have access to the SNMP agent residing in the system, at least one VLAN must have an IP address assigned to it.

For more information on assigning IP addresses, refer to Table 2-3.

SUPPORTED MIBS

Any Network Manager running SNMP can manage the BlackDiamond, provided the MIB is installed correctly on the management station. In addition to private MIBs, the BlackDiamond supports the standard MIBs listed in Table 2-6.

Table 2-6: Supported MIBs

Description	RFC Number
MIB II	1213
IP Forwarding Table MIB	1354
Bridge MIB	1493
Evolution of Interfaces	1573
RIP2 MIB	1724
RMON (Etherstats, History, Alarms, and Events)	1757
OSPF2 MIB	1850
RMON II Probe Configuration	2021
802.3 MAU MIB	2239

The IEEE Bridge MIB dot1dTpPortEntry PortInDiscards and dot1dBasePortEntry counters are not incremented.

CONFIGURING SNMP SETTINGS

The following SNMP parameters can be configured on the system:

- **Authorized trap receivers** — An authorized trap receiver can be one or more network management stations on your network. The BlackDiamond sends SNMP traps to all trap receivers. You can have a maximum of six trap receivers configured for each BlackDiamond. Entries in this list can be created, modified, and deleted using the RMON2 trapDestTable MIB variable, as described in RFC 2021.
- **Authorized managers** — An authorized manager can be either a single network management station, or a range of addresses (for example, a complete subnet) specified by a prefix and a mask. The BlackDiamond can have a maximum of 32 authorized managers.
- **Community strings** — The community strings allow a simple method of authentication between the system and the remote Network Manager. There are two types of community strings on the BlackDiamond. Read community strings provide read-only access to the system. The default read-only community string is *public*.

Read-write community strings provide read and write access to the system. The default read-write community string is *private.* A total of eight community strings can be configured on the system. The community string for all authorized trap receivers must be configured on the system for the trap receiver to receive system-generated traps.

- **System contact** (optional) — The system contact is a text field that enables you to enter the name of the person(s) responsible for managing the system.
- **System name** — The system name is the name that you have assigned to this system. The default name is the model name of the BlackDiamond (for example, BlackDiamond6800).
- **System location** (optional) — Using the system location field, you can enter an optional location for this system.

Table 2-7 describes SNMP configuration commands.

Table 2-7: SNMP Configuration Commands

Command	Description
enable snmp access	Turns on SNMP support for the BlackDiamond.
enable snmp trap	Turns on SNMP trap support.
config snmp add <ipaddress> {<mask>}	Adds the IP address of an SNMP management station to the access list. Up to 32 addresses can be specified.
config snmp add trapreceiver <ipaddress> community <string>	Adds the IP address of a specified trap receiver. The IP address can be a unicast, multicast, or broadcast. A maximum of six trap receivers is allowed.
config snmp community [readonly | readwrite] <string>	Adds an SNMP read or read/write community string. The default readonly community string is `public`. The default readwrite community string is `private`. Each community string can have a maximum of 127 characters.
config snmp delete [<ipaddress> {<mask>} | all]	Deletes the IP address of a specified SNMP management station or all SNMP management stations. If you delete all addresses, any machine can have SNMP management access to the switch.
config snmp delete trapreceiver [<ip_address> community <string> | all]	Deletes the IP address of a specified trap receiver or all authorized trap receivers.

Table 2-7: SNMP Configuration Commands (continued)

Command	Description
config snmp syscontact <string>	Configures the name of the system contact. A maximum of 255 characters is allowed.
config snmp sysname <string>	Configures the name of the system. A maximum of 255 characters is allowed. The default sysname is the model name of the device (for example, `BD6800`). The sysname appears in the BlackDiamond prompt.
config snmp syslocation <string>	Configures the location of the system. A maximum of 255 characters is allowed.

DISPLAYING SNMP SETTINGS

To display the SNMP settings configured on the BlackDiamond, enter the following command:

```
show management
```

This command displays the following information:

- Enable/disable state for telnet, SNMP, and Web access
- SNMP community strings
- Authorized SNMP station list
- SNMP trap receiver list
- Login statistics

RESETTING AND DISABLING SNMP

To reset and disable SNMP settings, use the commands in Table 2-8.

Table 2-8: SNMP Reset and Disable Commands

Command	Description
disable snmp access	Disables SNMP on the system.
disable snmp trap	Prevents SNMP traps from being sent from the system. Does not clear the SNMP trap receivers that have been configured.
unconfig management	Restores default values to all SNMP-related entries.

CHECKING BASIC CONNECTIVITY

The BlackDiamond offers the following commands for checking basic connectivity:

- `ping`
- `traceroute`

PING

The `ping` command enables you to send Internet Control Message Protocol (ICMP) echo messages to a remote IP device. The `ping` command is available for both the user and administrator privilege level.

The `ping` command syntax is

```
ping {continuous} {size <n>} <ip_address>
```

Options for the ping command are described in Table 2-9.

Table 2-9: Ping Command Parameters

Parameter	Description
`continuous`	Specifies ICMP echo messages to be sent continuously. This option can be interrupted by pressing any key.
`size <n>`	Specifies the size of the packet.

If a `ping` request fails, the system continues to send `ping` messages until interrupted. Press any key to interrupt a `ping` request.

TRACEROUTE

The `traceroute` command enables you to trace the routed path between the BlackDiamond system and a destination endstation. The `traceroute` command syntax is

```
traceroute <ip_address>
```

where `ip_address` is the IP address of the destination endstation.

3 Configuring Slots and Ports

Ports on the BlackDiamond can be configured in the following ways:

- Configuring a slot for a particular I/O module
- Enabling and disabling individual ports
- Configuring the port speed (Fast Ethernet ports only)
- Configuring half- or full-duplex mode
- Creating load-sharing groups on multiple ports
- Changing the Quality or Service (QoS) setting for individual ports

For more information on QoS, refer to Chapter 7.

CONFIGURING A SLOT

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated.

Once any of the ports on the module is configured (for example, a VLAN association, a VLAN tag configuration, or port parameters), all the port information and the module type for that slot must be saved to NVRAM. Otherwise, if the BlackDiamond is rebooted or the module is removed from the slot, the port, VLAN, and module configuration information is not saved.

For information on saving the configuration, refer to Chapter 13.

You can configure the BlackDiamond with the type of I/O module that is installed in each I/O slot. To do this, use the following command:

```
config slot <slot> module [f32t | f32f | g4x | g6x]
```

You can also pre-configure the slot before inserting the module card. This allows you to begin configuring the module and ports before installing the card in the chassis.

If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state, and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type. To clear the slot of a previously assigned I/O module type, use the following command:

```
clear slot <slot>
```

All configuration information related to the slot and the ports on the module is erased. If a module is present when you issue this command, the module is reset to default settings.

To display information about a particular slot, use the following command:

```
show slot <slot>
```

Information displayed includes the following:

- Card type, serial number, part number
- Current state (power down, operational, diagnostic, mismatch)
- Port information

SPECIFYING ONE OR MORE PORTS

On the BlackDiamond, the port number is a combination of the slot number and the port number. The nomenclature for the port number is as follows:

`slot:port`

For example, if a G4X I/O module (having a total of four ports) is installed in slot 2 of the BlackDiamond chassis, the following ports are valid:

- `2:1`
- `2:2`
- `2:3`
- `2:4`

You can also use wildcard combinations (*) to specify multiple BlackDiamond slot and port combinations. The following wildcard combinations are allowed:

- `slot:*` — Specifies all ports on a particular I/O module.
- `slot:`*`x`*`-slot:`*`y`* — Specifies a contiguous series of ports on a particular I/O module.
- `slot`*`a`*`:`*`x`* `- slot`*`b`*`:`*`y`* — Specifies a contiguous series of ports that begin on one I/O module and end on another I/O module.

ENABLING AND DISABLING PORTS

By default, all ports are enabled. To enable or disable one or more ports, use the following command:

`[enable | disable] port <slot:portlist>`

For example, to disable slot 7, ports 3, 5, and 12 through 15 on the BlackDiamond, enter the following:

`disable port 7:3,7:5,7:12-7:15`

Even though a port is disabled, the link remains enabled for diagnostic purposes.

CONFIGURING PORT SPEED AND DUPLEX SETTING

By default, the BlackDiamond is configured to use autonegotiation to determine the port speed and duplex setting for each port. You can select to manually configure the duplex setting and the speed of 10/100 Mbps ports, and you can manually configure the duplex setting on Gigabit Ethernet ports.

10BASE-T and 100BASE-TX ports can connect to either 10BASE-T or 100BASE-T networks. By default, the ports autonegotiate port speed. You can also configure each port for a particular speed (either 10 Mbps or 100 Mbps).

Gigabit Ethernet ports are statically set to 1 Gbps, and their speed cannot be modified.

To configure port speed and duplex setting, use the following command:

```
config port <slot:portlist> auto off {speed [10 | 100]} duplex [half |
full]
```

To configure the system to autonegotiate, use the following command:

```
config port <slot:portlist> auto on
```

TURNING OFF AUTONEGOTIATION FOR A GIGABIT ETHERNET PORT

In certain interoperability situations, you may need to turn autonegotiation off on a Gigabit Ethernet port. Even though a Gigabit Ethernet port runs only at full duplex, you must specify the duplex setting.

The following example turns autonegotiation off for port 1 on a G4X or G6X module located in slot 1:

```
config port 1:1 auto off duplex full
```

PORT COMMANDS

Table 3-1 describes the port commands.

Table 3-1: Port Commands

Command	Description
enable learning port <slot:portlist>	Enables MAC address learning on one or more ports. The default setting is enabled.
enable port <slot:portlist>	Enables a port.
enable sharing <slot:master_port> grouping <slot:portlist>	Defines a load-sharing group of ports. The ports specified in `<slot:portlist>` are grouped to the master port.
config port <slot:portlist> auto on	Enables autonegotiation for the particular port type; 802.3u for 10/100 Mbps ports or 802.3z for Gigabit Ethernet ports.
config port [<slot:portlist> | all] auto off {speed [10 | 100]} duplex [half | full]	Changes the configuration of a group of ports. Specify the following: ■ `auto off` — The port will not autonegotiate the settings ■ `speed` — The speed of the port (for 10/100 Mbps ports only) ■ `duplex` — The duplex setting (half- or full-duplex)
config port <slot:portlist> qosprofile <qosname>	Configures one or more ports to use a particular QoS profile.
config slot <slot> module [f32t | f32f | g4x | g6x]	Configures a slot for a particular I/O module card.
disable learning port <slot:portlist>	Disables MAC address learning on one or more ports for security purposes. If MAC address learning is disabled, only broadcast traffic and packets destined to a permanent MAC address matching that port number are forwarded. The default setting is enabled.
disable port <slot:portlist>	Disables a port. Even when disabled, the link is available for diagnostic purposes.
disable sharing <slot:master_port>	Disables a load-sharing group of ports.
clear slot <slot> module	Clears a slot of a previously assigned module type.
show port {<slot:portlist>} collisions	Displays real-time collision statistics.

Table 3-1: Port Commands (continued)

Command	Description
show port {<slot:portlist>} config	Displays the port configuration, including the following: ■ Port state ■ Link state ■ Link speed ■ Duplex mode ■ Flow control ■ Load-sharing information ■ Link media information
show port {<slot:portlist>} information	Displays detailed system-related information, including the following: ■ Port state ■ Link state ■ Autonegotiation state ■ Link speed ■ Duplex mode ■ Load-sharing information ■ Extreme Discovery Protocol (EDP) status ■ VLAN information ■ QoS information
show port {<slot:portlist>} packet	Displays a histogram of packet statistics.
show port {<slot:portlist>} qosmonitor	Displays real-time QoS statistics. For more information on QoS, refer to Chapter 7.
show port {<slot:portlist>} rxerrors	Displays real-time receive error statistics. For more information on error statistics, refer to Chapter 11.
show port {<slot:portlist>} stats	Displays real-time port statistics. For more information on port statistics, refer to Chapter 11.
show port {<slot:portlist>} txerrors	Displays real-time transmit error statistics. For more information on error statistics, refer to Chapter 11.
show port {<slot:portlist>} utilization	Displays real-time port utilization information. Use the [Spacebar] to toggle between packet, byte, and bandwidth utilization information.

Table 3-1: Port Commands (continued)

Command	Description
show slot <slot>	Displays slot-specific information, including the following: ■ Card type, serial number, part number ■ Current state ■ Port information

LOAD SHARING

Load sharing with BlackDiamond systems allows you to increase bandwidth and resilience by using a group of ports to carry traffic in parallel between switches. The sharing algorithm allows the BlackDiamond to use multiple ports as a single logical port. For example, VLANs see the load-sharing group as a single logical port. The algorithm also guarantees packet sequencing between clients.

If a port in a load-sharing group fails, traffic is redistributed to the remaining ports in the load-sharing group. If the failed port becomes active again, traffic is redistributed to include that port.

Load sharing is most useful in cases where the traffic transmitted from the BlackDiamond to the load-sharing group is sourced from an equal or greater number of ports on the system. For example, traffic transmitted to a 2-port load-sharing group should originate from a minimum of two other ports on the same system.

This feature is supported between BlackDiamond systems and Summit switches only, but may be compatible with third-party “trunking” or link-aggregation algorithms. Check with an Extreme Networks technical representative for more information.

CONFIGURING LOAD SHARING

To set up the BlackDiamond to load share among ports, you must create a load-sharing group of ports. Load-sharing groups are defined according to the following rules:

- Ports on the I/O module are divided into groups of two or four.
- Port in a load-sharing group must be part of the same I/O module.
- Ports in a load-sharing group must be contiguous.

- Follow the outlined boxes in Table 3-2, Table 3-3, and Table 3-4 to determine the valid port combinations.
- The first port in the load-sharing group is configured to be the “master” logical port. This is the reference port used in configuration commands. It can be thought of as the logical port representing the entire port group.

Table 3-2, Table 3-3, and Table 3-4 show the possible load-sharing port group combinations for the G4X module, the G6X module, and the F32T and F32F modules, respectively.

Table 3-2: Port Combinations for the G4X Module

Load-Sharing Group	1	2	3	4
4-port groups	x	x	x	x
2-port groups	x	x	x	x

Table 3-3: Port Combinations for the G6X Module

Load-Sharing Group	1	2	3	4	5	6
4-port groups	x	x	x	x		
2-port groups	x	x	x	x	x	x

Table 3-4: Port Combinations for the F32T and F32F Modules

Load-Sharing Group	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
4-port groups	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
2-port groups	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

Load-Sharing Group	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32
4-port groups	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
2-port groups	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

To define a load-sharing group, you assign a group of ports to a single, logical port number. To enable or disable a load-sharing group, use the following commands:

```
enable sharing <slot:master_port> grouping <slot:portlist>
disable sharing <slot:master_port>
```

The ports assigned to a load-sharing group cannot be members of a Spanning Tree Diamond (STPD).

The following example defines a load-sharing group on slot 3 that contains ports 9 through 12, and uses the first port in the group as the master logical port 9:

```
enable sharing 3:9 grouping 3:9-3:12
```

In this example, logical port 3:9 represents physical ports 3:9 through 3:12.

When using load sharing, you should always reference the master logical port of the load-sharing group (port 3:9 in the previous example) when configuring or viewing VLANs. VLANs configured to use other ports in the load-sharing group will have those ports deleted from the VLAN when load sharing becomes enabled.

VERIFYING THE LOAD-SHARING CONFIGURATION

The screen output resulting from the `show port config` command indicates the ports that are involved in load sharing and the master logical port identity.

PORT-MIRRORING

Port-mirroring configures the BlackDiamond to copy all traffic associated with one or more ports on a particular slot to a monitor port on the same slot. The monitor port can be connected to a network analyzer or RMON probe for packet analysis. The system uses a traffic filter that copies a group of traffic to the monitor port. The traffic filter can be defined based on one of the following criteria:

- MAC source address / destination address — All data sent to or received from a particular source or destination MAC address is copied to the monitor port.
- Physical port — All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port.

- VLAN — All data to and from a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- Virtual port — All data specific to a VLAN on a specific port is copied to the monitor port.

Up to 8 mirroring filters and one monitor port can be configured. Once a port is specified as a monitor port, it cannot be used for any other function.

Frames that contain errors are not mirrored.

PORT-MIRRORING COMMANDS

Port-mirroring commands are described in Table 3-5.

Table 3-5: Port-Mirroring Configuration Commands

Command	Description
enable mirroring port <slot:port>	Dedicates a port to be the mirror port.
config mirroring add [mac <mac_address> | vlan <name> | port <slot:port> | vlan <name> port <slot:port>]	Adds a single mirroring filter definition. Up to eight mirroring definitions can be added. You can mirror traffic from a MAC address, a VLAN, a physical port, or a specific VLAN/port combination.
config mirroring delete [mac <mac_address> | vlan <name> | port <slot:port> | vlan <name> port <slot:port> | all}	Deletes a particular mirroring filter definition, or all mirroring filter definitions.
disable mirroring	Disables port-mirroring.
show mirroring	Displays the port-mirroring configuration.

PORT-MIRRORING EXAMPLE

The following example selects port 3 on slot 7 as the mirror port, and sends all traffic coming into or out of the BlackDiamond on slot 7, port 1 to the mirror port:

```
enable mirroring port 7:3
config mirroring add port 7:1
```

The following example sends all traffic coming into or out of the system on slot 8, port 1 and the VLAN *default* to the mirror port:

```
config mirroring add port 8:1 vlan default
```

4 Virtual LANs (VLANs)

Setting up Virtual Local Area Networks (VLANs) on the BlackDiamond eases many time-consuming tasks of network administration while increasing efficiency in network operations.

This chapter describes the concept of VLANs and explains how to implement VLANs on the BlackDiamond.

Overview of Virtual LANs

The term “VLAN” is used to refer to a collection of devices that communicate as if they were on the same physical LAN. Any set of ports (including all ports on the switch) is considered a VLAN. LAN segments are not restricted by the hardware that physically connects them. The segments are defined by flexible user groups you create with the command-line interface.

Benefits

Implementing VLANs on your networks has the following advantages:

- **VLANs help to control traffic.**

  With traditional networks, congestion can be caused by broadcast traffic that is directed to all network devices, regardless of whether they require it. VLANs increase the efficiency of your network because each VLAN can be set up to contain only those devices that must communicate with each other.

- **VLANs provide extra security.**

  Devices within each VLAN can only communicate with member devices in the same VLAN. If a device in VLAN *Marketing* must communicate with devices in VLAN *Sales*, the traffic must cross a routing device.

- **VLANs ease the change and movement of devices.**

  With traditional networks, network administrators spend much of their time dealing with moves and changes. If users move to a different subnetwork, the addresses of each endstation must be updated manually.

  For example, with a VLAN, if an endstation in VLAN *Marketing* is moved to a port in another part of the network, and retains its original subnet membership; you must only specify that the new port is in VLAN *Marketing*.

TYPES OF VLANS

The BlackDiamond supports a maximum of 256 VLANs. VLANs can be created according to the following criteria:

- Physical port
- 802.1Q tag
- Ethernet, LLC SAP, or LLC/SNAP Ethernet protocol type
- A combination of these criteria

PORT-BASED VLANS

In a port-based VLAN, a VLAN name is given to a group of one or more ports on the switch. A port can be a member of only one port-based VLAN.

For example, on the G6X module in Figure 4-1, ports 1, 2, and 5 are part of VLAN *Marketing*; ports 3 and 4 are part of VLAN *Sales*; and port 6 is in VLAN *Finance*. On the F32T module, ports 1 through 4 and ports 17 through 20 are part of VLAN *Marketing*; ports 9 through 12 and ports 21 through 24 are part of VLAN *Sales*; and ports 9 through 16 and ports 25 through 32 are part of VLAN *Finance*.

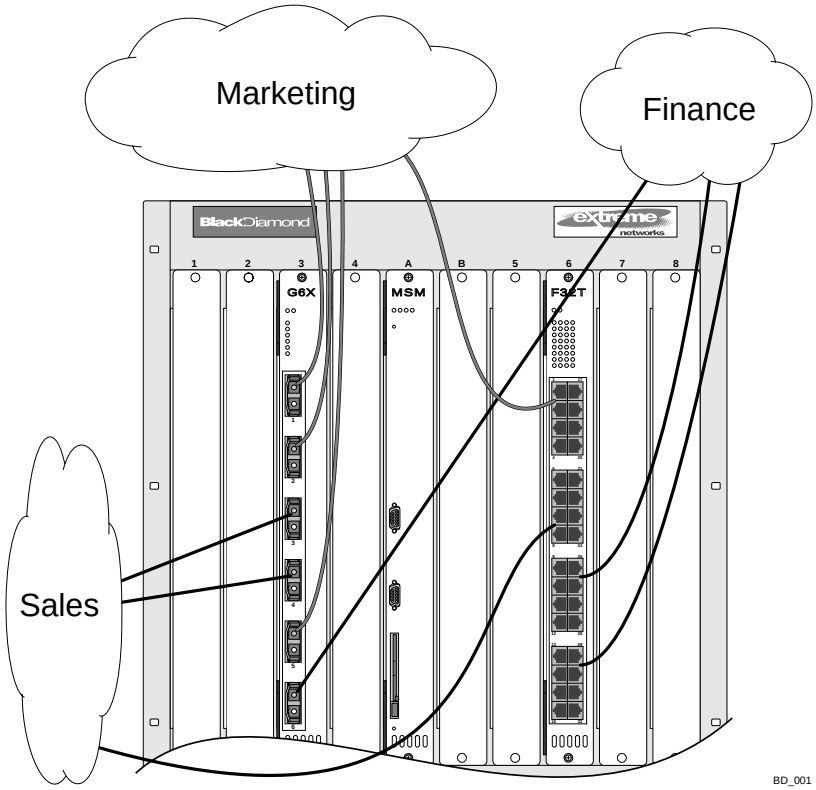

Figure 4-1: Example of a port-based VLAN

For the members of the different IP VLANs to communicate, the traffic must be routed by the BlackDiamond, even if they are physically part of the same I/O module. This means that each VLAN must be configured as a router interface with a unique IP address.

SPANNING SYSTEMS WITH PORT-BASED VLANS

To create a port-based VLAN that spans two systems, you must do two things:

- Assign the port on each system to the VLAN.
- Cable the two systems together using one port on each BlackDiamond per VLAN.

Figure 4-2 illustrates a single VLAN that spans two BlackDiamond systems. All ports on both systems belong to VLAN *Sales*. The two systems are connected using slot 1, port 4 on System 1, and slot 2, port 1 on System 2.

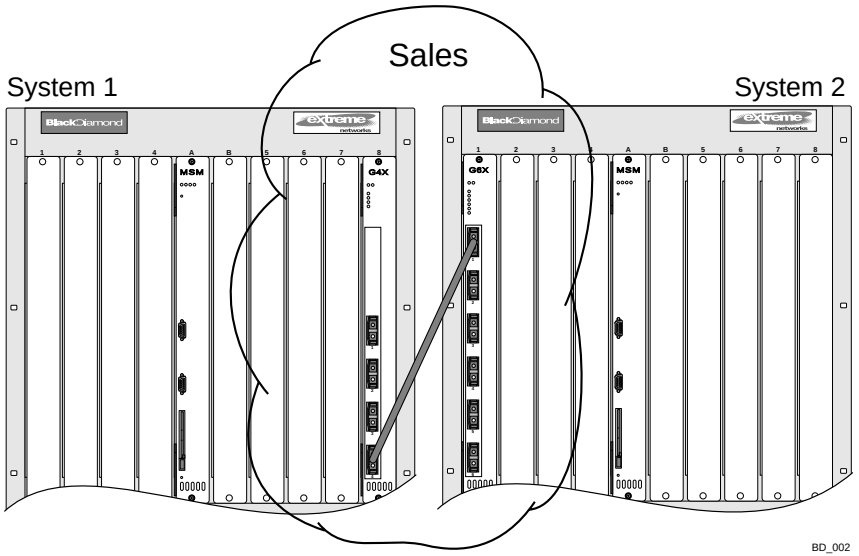

Figure 4-2: Single port-based VLAN spanning two switches

To create multiple VLANs that span two systems in a port-based VLAN, a port on System 1 must be cabled to a port on System 2 for each VLAN you want to have span across the systems. At least one port on each BlackDiamond must be a member of the corresponding VLANs, as well.

Figure 4-3 illustrates two VLANs spanning two switches. On System 1, all ports on slot 1 are part of VLAN *Accounting*; all port on slot 8 are part of VLAN *Engineering*. On System 2, all port on slot 3 are part of VLAN *Accounting*; all ports on slot 7 are part of VLAN *Engineering*.

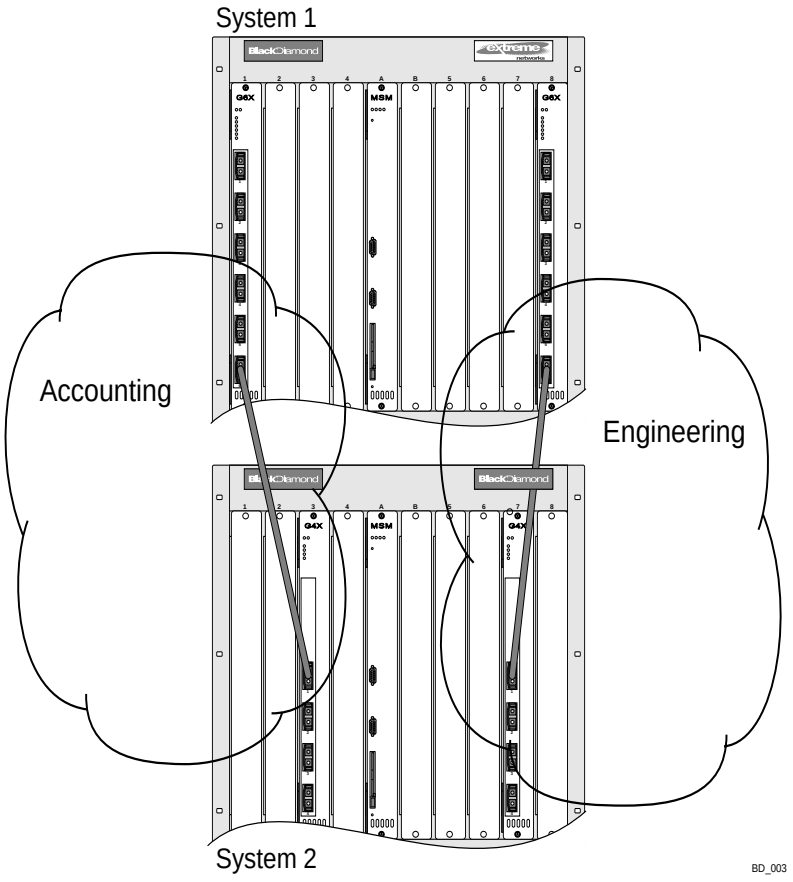

Figure 4-3: Two port-based VLANs spanning two BlackDiamond systems

VLAN *Accounting* spans System 1 and System 2 by way of a connection between Switch 1, slot 1, port 6 and Switch 2, slot 3, port 1. VLAN *Engineering* spans Switch 1 and Switch 2 by way of a connection between Switch 1, slot 8, port 6 and Switch 2, slot 7, port 1.

Using is configuration, you can create multiple VLANs that span multiple systems, in a daisy-chained fashion. Each BlackDiamond must have a dedicated port for each VLAN. Each dedicated port must be connected to a port that is a member of its VLAN on the next BlackDiamond.

TAGGED VLANS

Tagging is a process that inserts a marker (called a *tag*) into the Ethernet frame. The tag contains the identification number of a specific VLAN, called the *VLANid*.

The use of 802.1Q tagged packets may lead to the appearance of packets slightly bigger than the current IEEE 802.3/Ethernet maximum of 1518 bytes. This may affect packet error counters in other devices, and may also lead to connectivity problems if non-802.1Q bridges or routers are placed in the path.

USES OF TAGGED VLANS

Tagging is most commonly used to create VLANs that span systems. The system-to-system connections are typically called *trunks*. Using tags, multiple VLANs can span multiple systems using one or more trunks. In a port-based VLAN, each VLAN requires its own pair of trunk ports, as shown in Figure 4-3. Using tags, multiple VLANs can span two systems with a single trunk.

Another benefit of tagged VLANs is the ability to have a port be a member of multiple VLANs. This is particularly useful if you have a device (such as a server) that must belong to multiple VLANs. The device must have a NIC that supports 802.1Q tagging.

A single port can be a member of only one port-based VLAN. All additional VLAN membership for the port must be accompanied by tags. In addition to configuring the VLAN tag for the port, the server must have a *Network Interface Card (NIC)* that supports 802.1Q tagging.

ASSIGNING A VLAN TAG

Each VLAN may be assigned an 802.1Q VLAN tag. As ports are added to a VLAN with an 802.1Q tag defined, you decide whether each port will use tagging for that VLAN. The default mode of the BlackDiamond is to have all ports assigned to the VLAN named *default* with an 802.1Q VLAN tag (VLANid) of 1 assigned.

Not all ports in the VLAN must be tagged. As traffic from a port is forwarded out of the system, the BlackDiamond determines (in real time) if each destination port should use tagged or untagged packet formats for that VLAN. The BlackDiamond adds and strips tags, as required, by the port configuration for that VLAN.

Packets arriving tagged with a VLANid that is not configured in the BlackDiamond will be discarded.

Figure 4-4 illustrates the physical view of a network that uses tagged and untagged traffic.

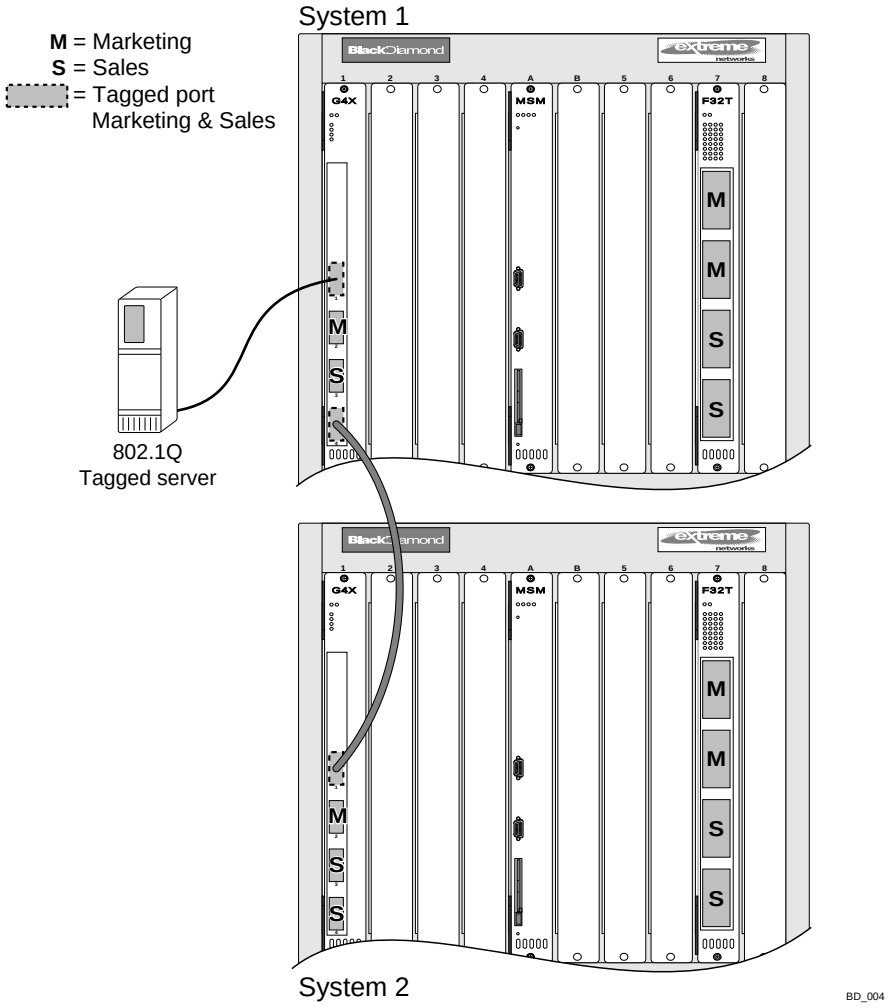

Figure 4-4: Physical diagram of tagged and untagged traffic

Figure 4-5 shows a logical diagram of the same network.

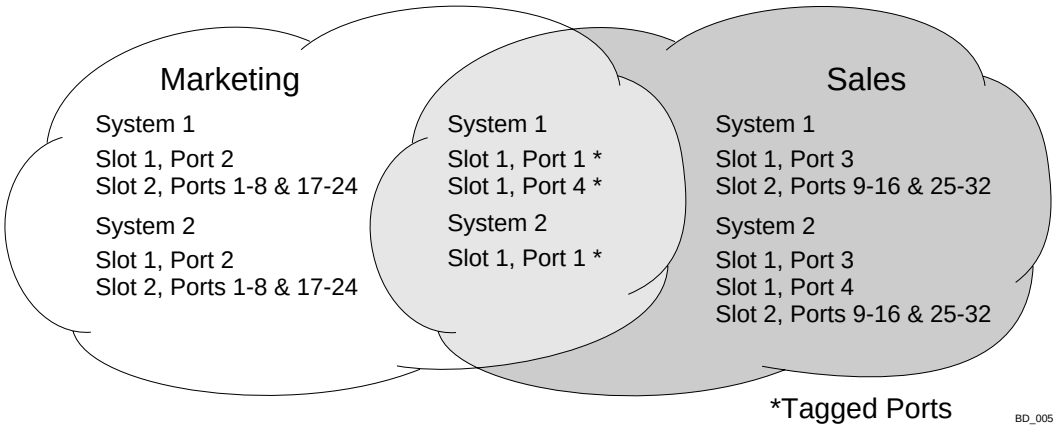

Figure 4-5: Logical diagram of tagged and untagged traffic

In Figure 4-4 and Figure 4-5:

- The trunk port on each system carries traffic for both VLAN *Marketing* and VLAN *Sales*.
- The trunk port on each switch is tagged.
- The server connected to slot 1, port 1 on System 1 has a NIC that supports 802.1Q tagging.
- The server connected to slot 1, port 1 on System 1 is a member of both VLAN *Marketing* and VLAN *Sales*.
- All other stations use untagged traffic.

As data passes out of the BlackDiamond, the system determines if the destination port requires the frames to be tagged or untagged. All traffic coming from and going to the server is tagged. Traffic coming from and going to the trunk ports is tagged. The traffic that comes from and goes to the other stations on this network is not tagged.

MIXING PORT-BASED AND TAGGED VLANS

You can configure the BlackDiamond using a combination of port-based and tagged VLANs. A given port can be a member of multiple VLANs, with the stipulation that only one of its VLANs uses untagged traffic. In other words, a port can simultaneously be a member of one port-based VLAN and multiple tag-based VLANs.

For the purposes of VLAN classification, packets arriving on a port with an 802.1Q tag containing a VLANid of zero are treated as untagged.

GENERIC VLAN REGISTRATION PROTOCOL

The Generic VLAN Registration Protocol (GVRP) allows a LAN device to signal other neighboring devices that it wishes to receive packets for one or more VLANs. The GVRP protocol is defined as part of the IEEE 802.1Q Virtual LANs draft standard. The main purpose of the protocol is to allow switches to automatically discover some of the VLAN information that would otherwise have to be manually configured in each switch. GVRP can also be run by network servers. These servers are usually configured to join several VLANs, and then signal the network switches of the VLANs of which they want to be part.

Figure 4-6 illustrates a network using GVRP.

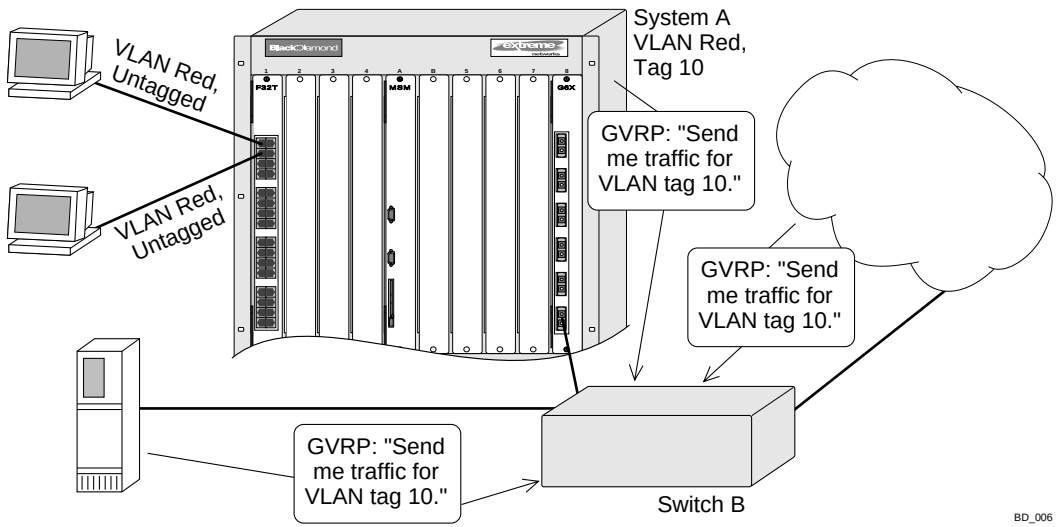

Figure 4-6: Network example using GVRP

In Figure 4-6, System A is a member of VLAN *Red*. VLAN *Red* has the VLANid 10. Slot 1, ports 1 and 2 on System A are added to the VLAN as untagged.

The configuration for System A is as follows:

```
create vlan red
config vlan red tag 10
config vlan red add port 1:1,1:2 untagged
enable gvrp
```

Switch B does not need to be configured with VLAN or tagging information. Instead, using GVRP, the server connected to Switch B, and the remainder of the network connected to Switch B provide Switch B with the information it needs to forward traffic. Switch A automatically adds slot 8, port 6 to VLAN *Red* because Switch A now knows that there are other devices on slot 8, port 6 that need access to VLAN *Red*.

VLANs that are automatically created using GVRP with the VLANid 10 are given names in the format

```
gvrp vlan xxxx
```

where *xxxx* is the VLANid (in decimal) that is discovered by GVRP. These VLANs are not permanently stored in nonvolatile RAM (NVRAM), and you cannot add or remove ports from these VLANs.

GVRP assumes that the VLANs for which it carries information operate using VLAN tags, unless explicitly configured otherwise. Typically, you must configure any untagged VLANs on the switches at the edges of the network, and the GVRP protocol is used across the core of the network to automatically configure other switches using tagged VLANs.

GVRP COMMANDS

Table 4-1 describes GVRP commands.

Table 4-1: GVRP Commands

Command	Description
enable gvrp	Enables the Generic VLAN Registration Protocol (GVRP). The default setting is disabled.
config gvrp {listen | send | both | none} {port <slot:portlist> | all}	Configures the sending and receiving GVRP information on one or more ports. Options include the following: ■ `listen` — Receive GVRP packets. ■ `send` — Send GVRP packets. ■ `both` — Send and receive GVRP packets. ■ `none` — Disable the port from participating in GVRP operation. The default setting is `both`.
disable gvrp	Disables the Generic VLAN Registration Protocol.
show gvrp	Displays the current configuration and status of GVRP.

PROTOCOL-BASED VLANS

Protocol-based VLANs enable you to define a packet filter that the BlackDiamond uses as the matching criteria to determine if a particular packet belongs to a particular VLAN.

Protocol-based VLANs are most often used in situations where network segments contain hosts running multiple protocols. For example, in Figure 4-7, the hosts are running both the IP and NetBIOS protocols.

The IP traffic has been divided into two IP subnets, 192.207.35.0 and 192.207.36.0. The subnets are internally routed by the BlackDiamond. The subnets are assigned different VLAN names, *Finance* and *Personnel*, respectively. The remainder of the traffic belongs to the VLAN named *MyCompany*. All ports are members of the VLAN *MyCompany*.

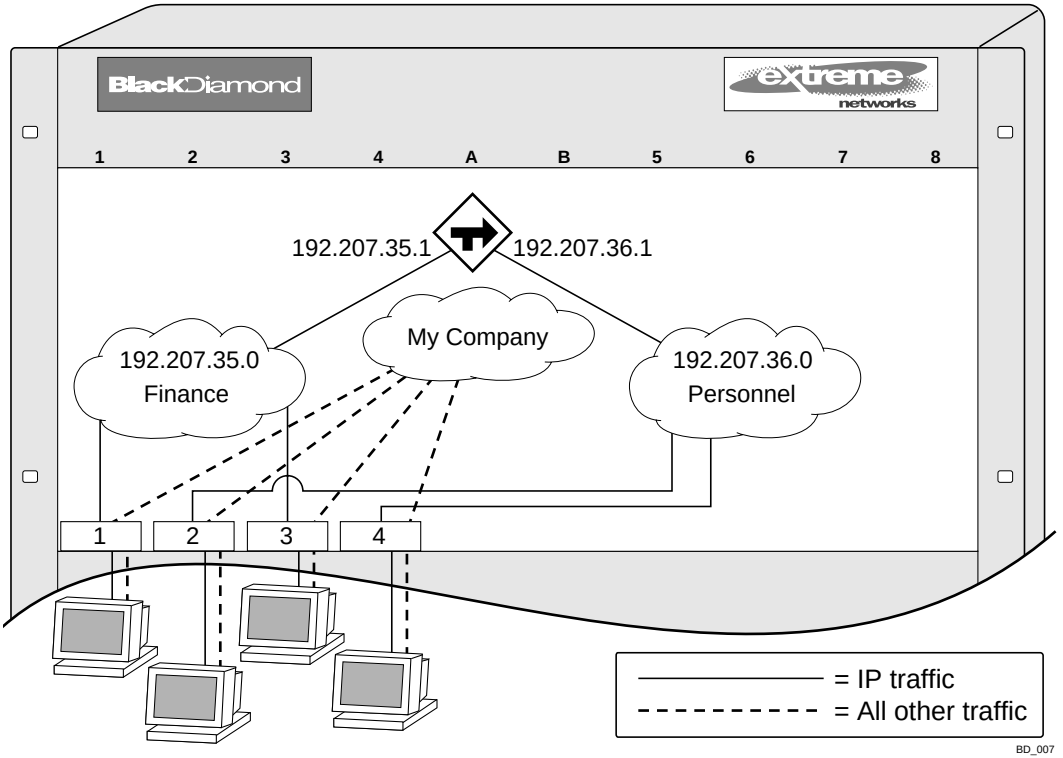

Figure 4-7: Protocol-based VLANs

PREDEFINED PROTOCOL FILTERS

The following protocol filters are predefined on the BlackDiamond:

- IP
- IPX
- NetBIOS
- DECNet
- IPX_8022
- IPX_SNAP
- AppleTalk

DEFINING PROTOCOL FILTERS

If necessary, you can define a customized protocol filter based on EtherType, LLC, and/or SNAP. Up to six protocols may be part of a protocol filter. To define a protocol filter, do the following:

1 Create a protocol using the following command:

`create protocol <protocol_name>`

For example:

`create protocol fred`

The protocol name can have a maximum of 31 characters.

2 Configure the protocol using the following command:

`config protocol <protocol_name> add <protocol_type> <hex_value>`

Supported protocol types include:

— `etype` — EtherType

The values for `etype` are four-digit hexadecimal numbers taken from a list maintained by the IEEE. This list can be found at the following URL:

`http://standards.ieee.org/regauth/ethertype/index.html`

— `llc` — LLC SAP

The values for `llc` are four-digit hexadecimal numbers that are created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP).

— `snap` — Ethertype inside an IEEE SNAP packet encapsulation.

The values for `snap` are the same as the values for `etype`, described previously.

For example:

`config protocol fred add llc feff`

`config protocol fred add snap 9999`

A maximum of seven protocol filters, each containing a maximum of six protocols, can be defined. However, no more than seven protocols should be active and configured for use.

For more information on SNAP protocol for Ethernet protocol types, see TR 11802-5:1997 (ISO/IEC) [ANSI/IEEE std. 802.1H, 1997 Edition].

DELETING A PROTOCOL FILTER

If a protocol filter is deleted from a VLAN, the VLAN is assigned a protocol filter of none. You can continue to configure the VLAN. However, no traffic is forwarded to the VLAN until a protocol is assigned to it.

PRECEDENCE OF TAGGED PACKETS OVER PROTOCOL FILTERS

If a VLAN is configured to accept tagged packets on a particular port, incoming packets that match the tag configuration take precedence over any protocol filters associated with the VLAN.

VLAN NAMES

The BlackDiamond supports up to 256 different VLANs. Each VLAN is given a name that can be up to 32 characters. VLAN names can use standard alphanumeric characters. The following characters are not permitted in a VLAN name:

- Space
- Comma
- Quotation mark

VLAN names must begin with an alphabetical letter. Quotation marks can be used to enclose a VLAN name that does not begin with an alphabetical character, or that contains a space, comma, or other special character.

VLAN names are locally significant. That is, VLAN names used on one system are only meaningful to that system. If another system is connected to it, the VLAN names have no significance to the other switch.

You should use VLAN names consistently across your entire network.

DEFAULT VLAN

The BlackDiamond ships with one default VLAN that has the following properties:

- The VLAN name is *default.*
- It contains all the ports on a new or initialized system.
- The default VLAN is untagged on all ports. It has an internal VLANid of 1.

CONFIGURING VLANS ON THE BLACKDIAMOND

This section describes the commands associated with setting up VLANs on the BlackDiamond. Configuring a VLAN involves the following steps:

1 Create and name the VLAN.

2 Assign an IP address and mask (if applicable) to the VLAN, if needed.

Each IP address and mask assigned to a VLAN must represent a unique IP subnet. You cannot configure the same IP subnet on different VLANs.

3 Assign a VLANid, if any ports in this VLAN will use a tag.

4 Assign one or more ports to the VLAN.

As you add each port to the VLAN, decide if the port will use an 802.1Q tag.

Table 4-2 describes the commands used to configure a VLAN.

Table 4-2: VLAN Configuration Commands

Command	Description
create vlan <name>	Creates a named VLAN.
create protocol <protocol_name>	Creates a user-defined protocol.
enable ignore-stp vlan <name>	Enables a VLAN from using STP port information. When enabled, all virtual ports associated with the VLAN are in STP forwarding mode. The default setting is disabled.
config dot1p ethertype <ethertype>	Configures an IEEE 802.1Q Ethertype. Use this command only if you have another switch that supports 802.1Q, but uses a different Ethertype value than 8100.

Table 4-2: VLAN Configuration Commands (continued)

Command	Description
config protocol <protocol_name> [add | delete] <protocol_type> <hex_value> {<protocol_type> <hex_value>} ...	Configures a protocol filter. Supported <protocol_type> values include: ■ etype ■ llc ■ snap The variable <hex_value> is a hexadecimal number between 0 and FFFF that represents either the Ethernet protocol type (for EtherType), the DSAP/SSAP combination (for LLC), or the SNAP-encoded Ethernet protocol type (for SNAP).
config vlan <name> ipaddress <ipaddress> {<mask>}	Assigns an IP address and an optional mask to the VLAN.
config vlan <name> [add | delete] port <slot:portlist> {tagged | untagged}	Adds one or more ports to a VLAN. You can specify tagged port(s), untagged port(s). By default, ports are untagged.
config vlan <name> delete port <slot:portlist> {tagged | untagged}	Deletes one or more ports from a VLAN.
config vlan <name> protocol [<protocol_name> | any]	Configures a protocol-based VLAN. If the keyword any is specified, then it becomes the default VLAN. All packets that cannot be classified into other protocol-based VLANs are assigned to the default VLAN of that port.
config vlan <name> qosprofile <qosname>	Configures a VLAN to use a particular QoS profile. Dynamic FDB entries associated with the VLAN are flushed once the change is committed.
config vlan <name> tag <vlanid>	Assigns a numerical VLANid. The valid range is from 1 to 4095.

VLAN CONFIGURATION EXAMPLES

The following example creates a port-based VLAN named *accounting*, assigns the IP address 132.15.121.1, and assigns slot 2, ports 1, 2, 3, and 6, and slot 4, ports 1 and 2 to it:

```
create vlan accounting
config accounting ipaddress 132.15.121.1
config default delete port 2:1-2:3,2:6,4:1,4:2
config accounting add port 2:1-2:3,2:6,4:1,4:2
```

Because VLAN names are unique, you do not need to enter the keyword `vlan` *after you have created the unique VLAN name. You can use the VLAN name alone.*

The following example creates a tag-based VLAN named *video*. It assigns the VLANid 1000. Slot 7, ports 4 through 8 are added as tagged ports to the VLAN.

```
create vlan video
config video tag 1000
config video add port 7:4-7:8 tagged
```

The following example creates a VLAN named *Sales*, with the VLANid 120. The VLAN uses both tagged and untagged ports. Slot 6, ports 1 through 3 are tagged, and slot 8, ports 16 though 18 are untagged. Note that, when not explicitly specified, ports are added as untagged.

```
create vlan sales
config sales tag 120
config sales add port 6:1-6:3 tagged
config sales add port 8:16-8:18
```

The following example creates a protocol-based VLAN named *IPSales*. Slot 5, ports 6 through 8, and slot 6, ports 1, 3, and 4-6 are assigned to the VLAN.

```
create vlan ipsales
config ipsales protocol ip
config ipsales add port 5:6-5:8,6:1,6:3-6:6
```

The following example defines a protocol filter, *myprotocol* and applies it to the VLAN named *myvlan*. This is an example only, and has no real-world application.

```
create protocol myprotocol
config protocol myprotocol add etype 0xf0f0
config protocol myprotocol add etype 0xffff
create vlan myvlan
config myvlan protocol myprotocol
```

DISPLAYING VLAN SETTINGS

To display VLAN settings, use the following command:

```
show vlan {<name> | all}
```

Sample output from this command is as follows:

```
show vlan all

VLAN Interface [1-fde] with name “net142” created by user
     Tagging:     Untagged (Internal tag 4095)
     IP:          Not configured. IGMP Snooping is disabled
     STPD:        Domain “s0” is running spanning tree protocol.
     Protocol     AppleTalk = SNAP:809b SNAP:80f3
     QoS Profile:         QP1
     Ports:       4.      (Number of active port=5)
                  Untag:     1:1  1:2  1:3  1:10

LAN Interface [2-fdd] with name “net123” created by user
     Tagging:     Untagged (Internal tag 4094)
     IP:          123.45.67.1/255.0.0.0
     STPD:        Domain “s0” is running spanning tree protocol.
     Protocol     Match all unfiltered protocols.
     QoS Profile:         QP1
     Ports:       15.      (Number of active port=1)
                  Untag:     2:6 2:7 2:8 2:9 2:10 2:11 2:12 2:13
                  2:14 2:15 6:2 6:3 6:4
```

The show command displays summary information about each VLAN, and includes the following:

- Name
- VLANid
- How the VLAN was created (manually or by GVRP)
- IP address
- STPD information
- Protocol information
- QOS profile information
- Ports assigned
- Tagged/untagged status for each port
- How the ports were added to the VLAN (manually or by GVRP)

To display protocol information, use the following command:

```
show protocol {<protocol> | all}
```

This show command displays protocol information, including the following:

- Protocol name
- List of protocol fields
- VLANs that use the protocol

DELETING VLANS

To delete a VLAN, or to return VLAN settings to their defaults, use the commands listed in Table 4-3.

Table 4-3: VLAN Delete and Reset Commands

Command	Description
disable ignore-stp vlan <name>	Allows a VLAN to use STP port information.
unconfig vlan <name> ipaddress	Resets the IP address of the VLAN.
delete vlan <name>	Removes a VLAN.
delete protocol <protocol>	Removes a protocol.

5 Forwarding Database (FDB)

This chapter describes the contents of the forwarding database (FDB), how the FDB works, and how to configure the FDB.

Overview of the FDB

The BlackDiamond maintains a database of all media access control (MAC) addresses received on all of its ports. It uses the information in this database to decide whether a frame should be forwarded or filtered.

FDB Contents

The database holds up to a maximum of 128K entries. Each entry consists of the MAC address of the device, an identifier for the port on which it was received, and an identifier for the VLAN to which the device belongs. Frames destined for devices that are not in the FDB are flooded to all members of the VLAN.

FDB Entry Types

The following are three types of entries in the FDB:

- **Dynamic entries** — Initially, all entries in the database are dynamic. Entries in the database are removed (aged-out) if, after a period of time (aging time), the device has not transmitted. This prevents the database from becoming full with obsolete entries by ensuring that when a device is removed from the network, its entry is deleted from the database. Dynamic entries are deleted from the database if the

BlackDiamond is reset or a power off/on cycle occurs. For more information about setting the aging time, refer to the section "Configuring FDB Entries," later in this chapter.

- **Non-aging entries** — If the aging time is set to zero, all aging entries in the database are defined as static, non-aging entries. This means that they do not age, but they are still deleted if the switch is reset.
- **Permanent entries** — Permanent entries are retained in the database if the BlackDiamond is reset or a power off/on cycle occurs. The system administrator must make entries permanent. A permanent entry can either be a unicast or multicast MAC address. All entries entered by way of the command-line interface are stored as permanent. The BlackDiamond can support a maximum of 64 permanent entries.

  Once created, permanent entries stay the same as when they were created. For example. the permanent entry store is not updated when any of the following take place:

  — A VLAN is deleted.

  — A VLANid is changed.

  — A port mode is changed (tagged/untagged).

  — A port is deleted from a VLAN.

  — A port is disabled.

  — A port enters blocking state.

  — A port QoS setting is changed.

  — A port goes down (link down).
- **Blackhole entrie**s — A blackhole entry configures packets with a specified MAC destination address to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination address must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle. Blackhole entries are never aged out of the database.

HOW FDB ENTRIES GET ADDED

Entries are added into the FDB in the following two ways:

- The BlackDiamond can learn entries. The system updates its FDB with the source MAC address from a packet, the VLAN, and the port identifier on which the source packet is received.
- You can enter and update entries using a MIB browser, an SNMP Network Manager, or the command-line interface.

ASSOCIATING A QOS PROFILE WITH AN FDB ENTRY

You can associate a QoS profile with a MAC address (and VLAN) of a device that will be dynamically learned. The FDB treats the entry like a dynamic entry (it is learned, it can be aged out of the database, and so on). The BlackDiamond applies the QoS profile as soon as the FDB entry is learned.

CONFIGURING FDB ENTRIES

To configure entries in the FDB, use the commands listed in Table 5-1.

Table 5-1: FDB Configuration Commands

Command	Description
create fdbentry <mac_address> vlan <name> [blackhole | <slot:portlist> | dynamic] {qosprofile <qosname>}	Creates an FDB entry. Specify the following: ■ `mac_address` — Device MAC address, using colon separated bytes. ■ `name` — VLAN associated with MAC address. ■ `blackhole` — Configures the MAC address as a blackhole entry. ■ `slot:portlist` — Slot and port numbers associated with MAC address. ■ `dynamic` — Specifies that the entry will be learned dynamically. Used to associated a QoS profile with a dynamically learned entry. ■ `qosname` — QoS profile associated with MAC address. If more than one port number is associated with a permanent MAC entry, packets are multicast to the multiple destinations.

Table 5-1: FDB Configuration Commands (continued)

Command	Description
config fdb agingtime <number>	Configures the FDB aging time. The range is 15 through 1,000,000 seconds. The default value is 300 seconds. A value of 0 indicates that the entry should never be aged out.
enable learning port <slot:portlist>	Enables MAC address learning on one or more ports.
disable learning port <slot:portlist>	Disables MAC address learning on one or more ports for security purposes. If MAC address learning is disabled, only broadcast traffic and packets destined to a permanent MAC address matching that port number, are forwarded. The default setting is enabled.

FDB CONFIGURATION EXAMPLES

This example adds a permanent entry to the FDB:

```
create fdbentry 00:E0:2B:12:34:56 vlan marketing port 3:4
```

The permanent entry has the following characteristics:

- MAC address is 00E02B123456.
- VLAN name is *marketing*.
- Slot number for this device is 3.
- Port number for this device is 4.

This example associates the QoS profile qp2 with a dynamic entry that will be learned by the FDB:

```
create fdbentry 00:A0:23:12:34:56 vlan net34 dynamic qosprofile qp2
```

This entry has the following characteristics:

- MAC address is 00A023123456.
- VLAN name is *net34*.
- The entry will be learned dynamically.
- QoS profile *qp2* will be applied when the entry is learned.

DISPLAYING FDB ENTRIES

To display FDB entries, use the command

```
show fdb {all | <mac_address> | vlan <name> | <slot:portlist> |
permanent | qos}
```

where the following is true:

- `all` — Displays all FDB entries.
- `mac_address` — Displays the entry for a particular MAC address.
- `vlan <name>` — Displays the entries for a VLAN.
- `slot:portlist` — Displays the entries for a slot and port combination.
- `permanent` — Displays all permanent entries.
- `qos` — Displays all entries that are associated with a QoS profile.

The following sample output shows the information displayed when you request output for all FDB entries:

```
show fdb

Index      Mac            Vlan        Flags PortList
-------------------------------------------------------------------
02b7: 0 00:e0:2b:00:57:00 Default(0001) s m CPU
02b9: 0 00:00:81:05:00:aa Default(0001) d m 2:1
0d2d: 0 00:10:4b:07:35:a4 Default(0001) d m 2:1
03bd: 0 00:60:08:18:a9:bd Default(0001) d m 2:1
0f00: 0 ff:ff:ff:ff:ff:ff Default(0001) s m CPU,2:1,6:3
1237: 0 00:a0:c9:20:1d:92 Default(0001) d m 2:1
1288: 0 00:e0:2b:00:00:00        (0000) s m CPU
1289: 0 00:e0:2b:00:00:00 Default(0001) d m 2:1
12f5: 0 00:e0:2b:00:ae:00 Default(0001) d m 6:3
12f6: 0 00:e0:2b:00:a8:00 Default(0001) d m 2:1
1636: 0 00:a0:c9:20:1C:a2 Default(0001) d m 2:1
173d: 0 01:80:c2:00:00:00        (0000) s m CPU
17eb: 0 00:a0:c9:6c:cf:55 Default(0001) d m 2:1
1892: 0 01:00:5e:00:00:02 Default(0001) s m CPU,2:1,6:3
19d2: 0 01:00:5e:00:00:01 Default(0001) s m CPU,2:1,6:3

Total: 16 Static: 7 Perm: 0 Dyn: 9 Dropped: 0
FDB Aging time: 300 seconds
```

The show command displays summary information, including

- MAC address
- VLAN name and VLANID

  The VLANID 0000 indicates that the entry is a special entry that is not associated with any one VLAN.
- Entry method (shown in the field labeled *Flags*):

  s — Static entry configured by the user

  d — Dynamic entry learned by the switch

  m — MAC address entry

  i — MAC address entry that is used for IP routing
- Slot and port

REMOVING FDB ENTRIES

You can remove one or more specific entries from the FDB, or you can clear the entire FDB of all entries by using the commands listed in Table 5-2.

Table 5-2: Removing FDB Entry Commands

Command	Description
delete fdbentry <mac_address> vlan <name>	Deletes a permanent FDB entry.
clear fdb [all | <mac_address> | vlan <name> | <slot:portlist>]	Clears dynamic FDB entries that match the filter. Use the keyword all to clear all dynamic entries.

6 Spanning Tree Protocol (STP)

Using the Spanning Tree Protocol (STP) functionality of the BlackDiamond makes your network more fault tolerant. The following sections explain more about STP and the STP features supported by the system.

STP is a part of the 802.1D bridge specification defined by the IEEE Computer Society. To explain STP in terms used by the 802.1D specification, the BlackDiamond will be referred to as a bridge.

Overview of the Spanning Tree Protocol

STP is a bridge-based mechanism for providing fault tolerance on networks. STP allows you to implement parallel paths for network traffic, and ensure that

- Redundant paths are disabled when the main paths are operational.
- Redundant paths are enabled if the main path fails.

Spanning Tree Domains

The BlackDiamond can be partitioned into multiple virtual bridges. Each virtual bridge can run an independent Spanning Tree instance. Each Spanning Tree instance is called a *Spanning Tree Domain* (STPD). Each STPD has its own Root Bridge and active path. Once the STPD is created, one or more VLANs can be assigned to it.

A port can belong to only one STPD. If a port is a member of multiple VLANs, then all those VLANs must belong to the same STPD.

The key points to remember when configuring VLANs and STP are the following:

- Each VLAN forms an independent broadcast domain.
- STP blocks paths to create a loop-free environment.
- When STP blocks a path, no data can be transmitted or received on the blocked port.
- Within any given STPD, all VLANs belonging to it use the same spanning tree.

Care must be taken to ensure that STPD instances within a single BlackDiamond do not see each other in the same broadcast domain. This could happen if, for example, another external bridge is used to connect VLANs belonging to separate STPDs.

If you delete a STPD, the VLANs that were members of that STPD are also deleted. You must remove all VLANs associated with the STP before deleting the STPD.

DEFAULTS

The default device configuration contains a single STPD called *s0*. The default VLAN is a member of STPD s0.

All STP parameters default to the IEEE 802.1D values, as appropriate.

STP CONFIGURATIONS

When you assign VLANs to an STPD, pay careful attention to the STP configuration and its effect on the forwarding of VLAN traffic.

Figure 6-1 illustrates a network that uses VLAN tagging for trunk connections. The following four VLANs have been defined:

- *Sales* is defined on Switch A, Switch B, and Switch M.
- *Personnel* is defined on Switch A, Switch B, and Switch M.
- *Manufacturing* is defined on Switch Y, Switch Z, and Switch M.
- *Engineering* is defined on Switch Y, Switch Z, and Switch M.
- *Marketing* is defined on all switches (Switch A, Switch B, Switch Y, Switch Z, and Switch M).

Two STPDs are defined:

- STPD1 contains VLANs *Sales* and *Personnel.*
- STPD2 contains VLANs *Manufacturing* and *Engineering.*

The VLAN *Marketing* is a member of the default STPD, but not assigned to either STPD1 or STPD2.

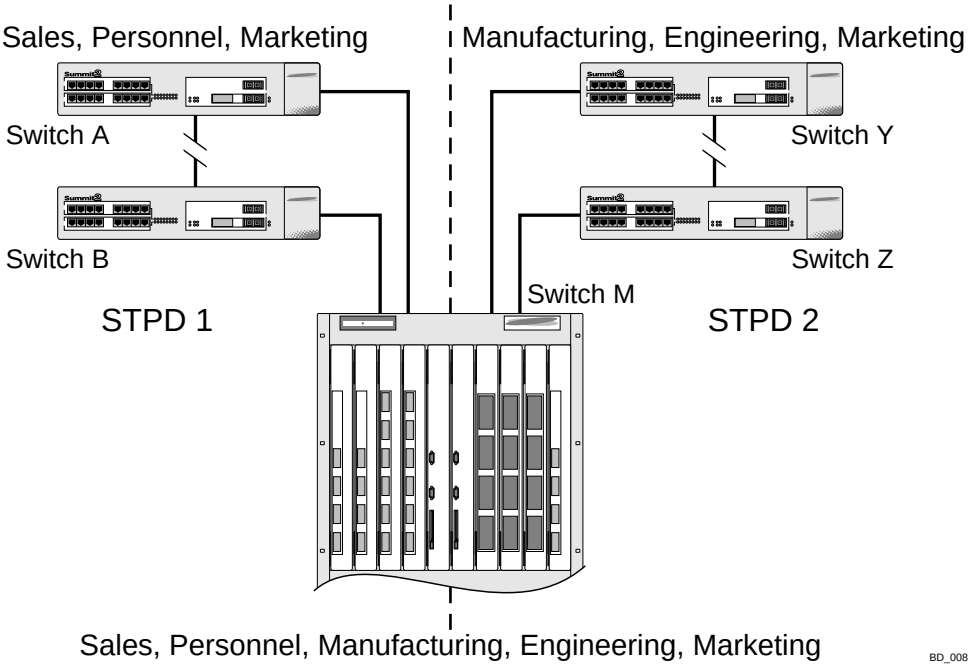

Figure 6-1: Multiple Spanning Tree Domains

When the switches in this configuration start up, STP configures each STP domain such that there are no active loops in the topology. STP could configure the topology in a number of ways to make it loop-free.

In Figure 6-1, the connection between Switch A and Switch B is put into blocking state, and the connection between Switch Y and Switch Z is put into blocking state. After STP converges, all the VLANs can communicate, and all bridging loops are prevented.

The VLAN *Marketing*, which has not been assigned to either STPD1 or STPD2, communicates using all five switches. The topology has no loops, because STP has already blocked the port connection between Switch A and Switch B, and between Switch Y and Switch Z.

Within a single STPD, you must be extra careful when configuring your VLANs. Figure 6-2 illustrates a network that has been incorrectly set up using a single STPD so that the STP configuration disables the ability of the switches to forward VLAN traffic.

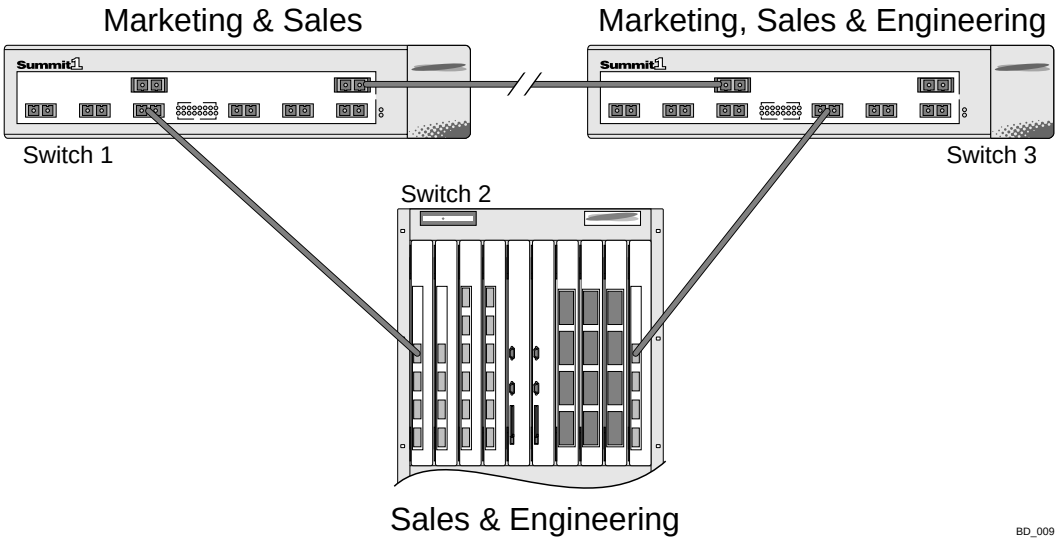

Figure 6-2: Tag-based STP configuration

The tag-based network in Figure 6-2 has the following configuration:

- Switch 1 contains VLAN *Marketing* and VLAN *Sales*.
- Switch 2 contains VLAN *Engineering* and VLAN *Sales*.
- Switch 3 contains VLAN *Marketing*, VLAN *Engineering*, and VLAN *Sales*.
- The tagged trunk connections for three switches form a triangular loop that is not permitted in an STP topology.
- All VLANs in each switch are members of the same STPD.

STP may block traffic between Switch 1 and Switch 3 by disabling the trunk ports for that connection on each switch.

Switch 2 has no ports assigned to VLAN marketing. Therefore, if the trunk for VLAN marketing on Switches 1 and 3 is blocked, the traffic for VLAN marketing will not be able to traverse the switches.

CONFIGURING STP ON THE BLACKDIAMOND

STP configuration involves the following actions:

- Create one or more STP domains using the following command:

  ```
  create stpd <stpd_name>
  ```

  STPD, VLAN, and QoS profile names must all be unique. For example, a name used to identify a VLAN cannot be used when you create an STPD or a QoS profile.

- Add one or more VLANs to the STPD using the following command:

  ```
  config stpd <stpd_name> add vlan <name>
  ```

  The ports assigned to a STPD cannot be members of a load-sharing group.

- Enable STP for one or more STP domains using the following command:

  ```
  enable stpd [<stpd_name> | all]
  ```

  All VLANs belong to a STPD. If you do not want to run STP on a VLAN, you must add the VLAN to a STPD that is disabled.

Once you have created the STPD, you can optionally configure STP parameters for the STPD.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The following parameters can be configured on each STPD:

- Hello time
- Forward delay
- Max age
- Bridge priority

The following parameters can be configured on each port:

- Path cost
- Port priority

The device supports the RFC 1493 Bridge MIB. Parameters of only the s0 default STPD are accessible through this MIB.

Table 6-1 shows the commands used to configure STP.

Table 6-1: STP Configuration Commands

Command	Description
create stpd <stpd_name>	Creates an STPD. When created, an STPD has the following default parameters: ■ Bridge priority — 32,768 ■ Hello time — 2 seconds ■ Forward delay — 15 seconds
enable stpd [<stpd_name> | all]	Enables the STP protocol for one or all STPDs. The default setting is disabled.
enable stpd port <slot:portlist>	Enables the STP protocol on one or more ports. If STPD is enabled for a port, Bridge Protocol Data Units (BPDUs) will be generated on that port if STP is enabled for the associated STPD. The default setting is enabled.
config stpd <stpd_name> add vlan <name>	Adds a VLAN to the STPD.
config stpd <stpd_name> hellotime <value>	Specifies the time delay (in seconds) between the transmission of BPDUs from this STPD when it is the Root Bridge. The range is 1 through 10. The default setting is 2 seconds.
config stpd <stpd_name> forwarddelay <value>	Specifies the time (in seconds) that the ports in this STPD spend in the listening and learning states when the switch is the Root Bridge. The range is 4 through 30. The default setting is 15 seconds.

Table 6-1: STP Configuration Commands (continued)

Command	Description
config stpd <stpd_name> maxage <value>	Specifies the maximum age of a BPDU in this STPD. The range is 6 through 40. The default setting is 20 seconds. Note that the time must be greater than, or equal to 2 * (Hello Time + 1) and less than, or equal to 2 * (Forward Delay –1).
config stpd <stpd_name> priority <value>	Specifies the priority of the STPD. By changing the priority of the STPD, you can make it more or less likely to become the Root Bridge. The range is 0 through 65,535. The default setting is 32,768. A setting of 0 indicates the highest priority.
config stpd <stpd_name> port cost <value> <slot:portlist>	Specifies the path cost of the port in this STPD. The range is 1 through 65,535. The switch automatically assigns a default path cost based on the speed of the port, as follows: ■ For a 10Mbps port, the default cost is 100. ■ For a 100Mbps port, the default cost is 19. ■ For a 1000Mbps port, the default cost is 4.
config stpd <stpd_name> port priority <value> <slot:portlist>	Specifies the priority of the port in this STPD. By changing the priority of the port, you can make it more or less likely to become the Root Port. The range is 0 through 255. The default setting is 128. A setting of 0 indicates the lowest priority.

CONFIGURATION EXAMPLE

The following example creates and enables an STPD named *Backbone_st*. It assigns the *Manufacturing* VLAN to the STPD. It disables STP on slot 2, ports 1 through 7, and slot 3 port 12.

```
create stpd backbone_st
config stpd backbone_st add vlan manufacturing
enable stpd backbone_st
disable stpd backbone_st port 2:1-2:7,3:12
```

DISPLAYING STP SETTINGS

To display STP settings for all ports, use the following command:

```
show stpd {<stpd_name> | all}
```

This command displays the following information:

- STPD name
- Bridge ID
- STPD configuration information

Sample output from the command is as follows:

```
show stpd s0

Stpd:s0                Stp:DISABLED        Number of Ports:13
Ports: 2:16,2:17,2:18,2:19,2:20,2:21,2:22,2:23,2:24,2:25,6:1,6:2,6:4
Vlans:   Default
BridgeID           80:00:00:e0:2b:00:57:00
Designated root:   00:00:00:00:00:00:00:00
RootPathCost: 0
MaxAge: 0s         HelloTime: 0s       ForwardDelay: 0s
CfgBrMaxAge: 20s   CfgBrHelloTime: 2s  CfgBrForwardDelay:15s
Topology Change Time: 35s              Hold time: 1s
Topology Change Detected: FALSE        Topology Change:FALSE
Number of Topology Changes: 0
Time Since Last Topology Change: 0s
```

To display the STP state of a port, use the following command:

```
show stpd <stpd_name> port <slot:portlist>
```

This command displays the following:

- STPD port configuration
- STPD state (root bridge, and so on)
- STPD port state (forwarding, blocking, and so on)

DISABLING AND RESETTING STP

To disable STP or return STP settings to their defaults, use the commands listed in Table 6-2.

Table 6-2: STP Disable and Reset Commands

Command	Description
delete stpd <stpd_name>	Removes an STPD. An STPD can only be removed if all VLANs have been deleted from it.
disable stpd [<stpd_name> | all]	Disables the STP mechanism on a particular STPD, or for all STPDs.
disable stpd port <slot:portlist>	Disables STP on one or more ports. Disabling STP on one or more ports puts those ports in FORWARDING state; all BPDUs received on those ports will be disregarded.
unconfig stpd {<stpd_name> | all}	Restores default STP values to a particular STPD or to all STPDs.

7 Quality of Service (QoS)

This chapter describes the concept of Quality of Service (QoS) and explains how to implement QoS on the BlackDiamond.

Overview of Quality of Service

QoS is a feature of the BlackDiamond that allows you to specify different service levels for outbound traffic. QoS is an effective control mechanism for networks that have heterogeneous traffic patterns. Using QoS, you can specify the service that a traffic type receives.

The main benefit of QoS is that it allows you to have control over the types of traffic that receive enhanced service from the system. For example, if video traffic requires a higher priority than data traffic, using QoS you can assign a different QoS profile to those VLANs that are transmitting video traffic.

Building Blocks

QoS is determined by one or more of the following building blocks:

- **QoS mode** — Indicates whether the system should use egress or ingress traffic classifications. Ingress is the default.
- **QoS profile** — Includes bandwidth and priority parameters.
- **Traffic classification** —A method of grouping traffic that has one or more attributes in common.

QoS profiles are assigned to traffic classifications, independent of the QoS mode chosen, in order to modify switch forwarding behavior.

QOS MODE

There are two modes of QoS. Ingress mode, the default, can use a wide variety of traffic classifications, but has a limitation of being able to use only the default four QoS profiles. You can modify the bandwidth parameters of the default QoS profiles.

Using egress mode, you can define additional QoS profiles, but you must use a smaller selection of traffic classifications. You can modify both the prioritization and bandwidth parameters of user-defined QoS profiles.

DEFAULT QOS PROFILES

Four default QoS profiles are provided that cannot be deleted. The default QoS profile names are as follows:

- qp1
- qp2
- qp3
- qp4

The default QoS profiles exist in either ingress or egress mode. In ingress mode, only the default QoS profiles are observed. In egress mode, up to 28 additional custom profiles may be defined, for a total of 32. You cannot create custom profiles in ingress mode.

The parameters that make up a QoS profile include the following:

- **Minimum bandwidth** — The minimum percentage of bandwidth that the traffic requires. The system is required to provide the minimum amount of bandwidth to the traffic. The lowest possible value is 0%.
- **Maximum bandwidth** — The maximum percentage of bandwidth that the traffic is permitted to use.

- **Priority** — The level of priority in which the traffic will be serviced by the switch. Choices include:
  - Low
  - Normal
  - Medium
  - High

A QoS profile does not alter the behavior of the BlackDiamond until it is assigned to a traffic classification.

The details of the default profiles are shown in Table 7-1.

Table 7-1: Default QoS Profiles

Profile Name	Priority	Minimum Bandwidth	Maximum Bandwidth
qp1	Low	0%	100%
qp2	Normal	0%	100%
qp3	Medium	0%	100%
qp4	High	0%	100%

You can modify the minimum and maximum bandwidth parameters of the default QoS profiles in either ingress or egress mode. The priority parameter can not be modified in ingress mode.

TRAFFIC GROUPINGS

Different traffic groupings are available, depending on the QoS mode configured for the system. In the event that a given packet matches two or more grouping criteria, there is a predetermined precedence for which traffic grouping will apply. By default, all traffic groupings are placed in the QoS profile named *qp1*.

The available traffic groupings are listed in the following sections, in order of highest to lowest precedence.

INGRESS TRAFFIC GROUPINGS

Ingress traffic groupings include the following:

- IP destination address — A specific QoS profile can be associated with an IP destination address, or range of IP destination addresses specified using a subnet mask. The QoS parameters are dynamically associated with a route when the route table is built. This is controlled by the following command:

  ```
  config ipqos [add | delete] <ip_destination_address>/<mask_length>
  qosprofile <qosname>
  ```

- Destination MAC address — When making a permanent FDB entry, you can provide a QoS profile. You can also provide a QoS profile that will be bound to a dynamic FDB entry when the MAC address is learned. This is configured using the following command:

  ```
  create fdbentry <mac addr> vlan <vlan name> [blackhole | port
  <slot:portlist> | dynamic] qosprofile <qosname>
  ```

  For example:

  `create fdbentry 00:11:22:33:44:55 vlan default port 4:1 qosprofile qp1`

- IEEE 802.1p — When traffic that contains 802.1p prioritization bits are seen, the traffic is mapped to the four default QoS profiles. No user configuration is required for this type of traffic grouping. Table 7-2 describes 802.1p values and their associated QoS profiles.

Table 7-2: 802.1p Values and Associated QoS Profiles

802.1p Value	QoS Profile
0	qp1
1	qp1
2	qp2
3	qp2
4	qp3
5	qp3
6	qp4
7	qp4

- PACE — When 3Com PACE traffic is seen, it is mapped to the profile named *qp3*. Observance of PACE can be controlled by using the following command:

  ```
  {enable | disable} pace
  ```

- Source port — You can configure a QoS profile to all the traffic being received from a particular port. This is controlled by using the following command:

  ```
  config port <port:portlist> qosprofile <qosname>
  ```

- VLAN — This is controlled by using the following command:

  ```
  config vlan <name> qosprofile <qosname>
  ```

EGRESS TRAFFIC GROUPINGS

Egress traffic groupings include the following:

- IP destination address — A specific QoS profile can be associated with an IP destination address, or range of IP destination addresses specified using a subnet mask. The QoS parameters are dynamically associated with a route when the route table is built. This is controlled by the following command:

  ```
  config ipqos [add | delete] <destination_address> qosprofile
  <qosname>
  ```

- Destination MAC address — Configuration is as described in the section, "Ingress Traffic Groupings," except that bandwidth parameters of the QoS profile are observed.
- VLAN — Configuration is as described in the section, "Ingress Traffic Groupings," except that bandwidth parameters of the QoS profile are observed.

PRECEDENCE

If traffic falls into multiple traffic groupings, the following order of precedence applies:

Ingress mode

- IP destination address
- Destination MAC address
- 802.1p prioritization bits
- PACE
- Source port
- VLAN

Egress mode

- IP destination address
- Destination MAC address
- VLAN

PRIORITIZATION

Prioritization is used when there is bandwidth contention for transmission on a port. The four levels of priority are used as a mechanism for resolving the contention between traffic groups. If traffic groups have the same priority, a "round-robin" algorithm is applied.

CREATING AND CONFIGURING A QOS PROFILE

Up to 28 custom QoS profiles can be created on the BlackDiamond in egress mode. To create a QoS profile, use the following command:

```
create qosprofile <name>
```

A new QoS profile is created with the following default values:

- Minimum bandwidth — 0%
- Maximum bandwidth — 100%
- Priority — low

Each of the default values is configurable by using the following command in egress mode:

```
config qosprofile <qosname> {minbw <percent>} {maxbw <percent>}
{priority <level>}
```

In ingress mode, use the same command syntax to modify the bandwidth parameters of the default QoS profiles. However, the priority level specified is ignored.

ASSIGNING A QOS PROFILE

Once you have established one or more traffic classifications and configured one or more QoS profiles, you can match them together using one of the following commands:

```
config vlan <name> qosprofile <qosname>
```

or

```
config port <slot:portlist> qosprofile <qosname>
```

You can assign a QoS profile to a MAC entry by using the following command:

```
create fdbentry <mac_address> vlan <name> [blackhole | <slot:portlist>
| dynamic] {qosprofile <qosname>}
```

You can assign a QoS profile to an IP address by using the following command:

```
config ipqos [add | delete] <ip_destination_address> qosprofile
<qosname>
```

PORT QUEUE MONITOR

There are multiple queues per physical port in the BlackDiamond. In ingress mode, the four QoS profiles align to the four queues (for example, qp1 to the first queue, qp2 to the second queue, and so on). In egress mode, the mapping of QoS profiles to queues is a function of the particular configuration.

The Port Queue Monitor (PQM) is a utility that monitors all the queues assigned to a port. The PQM monitors the number of frames and the frames per second a specific queue is responsible for transmitting on a physical port. The real-time display scrolls through the given portlist to provide these statistics. The particular port being monitored at that time is indicated by an asterisk (*) appearing after the port number in the display.

When the real-time display is not used, the PQM monitors the queue usage in the background. Error messages are generated by the PQM if the traffic exceeds the parameters of the QoS profile(s).

Table 7-3 describes the PQM commands.

Table 7-3: PQM Commands

Command	Description
enable qosmonitor {port <slot:portlist>}	Enables the QoS monitoring capability. When no port is specified, the QoS monitor automatically samples all the ports and records the sampled results. Error messages are logged to the syslog if the traffic exceeds the parameters of the QoS profile(s). The default setting is disabled.
disable qosmonitor	Disables the QoS monitoring capability.
show port {<slot:portlist>} qosmonitor	Displays real-time QoS statistics for one or more ports.

CONFIGURING QOS

Table 7-4 describes the commands used to configure QoS.

Table 7-4: QoS Configuration Commands

Command	Description
enable pace	Enables recognition of the PACE bit. Available in ingress mode, only.
create qosprofile <qosname>	Creates a QoS profile. The default values assigned to a created QoS profile are ■ Minimum bandwidth — 0% ■ Maximum bandwidth — 100% ■ Priority — low
config qosmode [ingress | egress]	Changes the QoS mode to ingress mode or egress mode.

Table 7-4: QoS Configuration Commands (continued)

Command	Description
config qosprofile <qosname> {minbw <percent>} {maxbw <percent>} {priority <level>}	Configures a QoS profile. Specify: ■ `minbw` — The minimum bandwidth percentage guaranteed to be available to this queue. The default setting is 0. ■ `maxbw` — The maximum bandwidth percentage this queue is permitted to use. The default setting is 100. ■ `priority` — The service priority for this queue. Settings include low, medium-low, medium, high. The default setting is low. Available in egress mode, only.
config port <slot:portlist> qosprofile <qosname>	Allows you to configure one or more ports to use a particular QoS profile. Available in ingress mode, only.
config vlan <name> qosprofile <qosname>	Allows you to configure a VLAN to use a particular QoS profile.
disable pace	Disables recognition of the PACE bit. Available in ingress mode, only.

SAMPLE INGRESS MODE QOS CONFIGURATION

The following ingress mode example modifies an existing QoS profile and applies it to a VLAN traffic grouping. The priority parameter, although required, is ignored when configuring a default QoS profile in ingress mode.

```
config qosprofile qp4 minbw 15% maxbw 100% priority high
config vlan sales qosprofile qp4
```

SAMPLE EGRESS MODE QOS CONFIGURATION

This egress mode example does the following:

- Configures the QoS mode of the switch.
- Creates a QoS profile *mktgqos*, with the following characteristics:
  - minimum bandwidth = 0%
  - maximum bandwidth = 10%
  - priority = low
- Applies the QoS profile *mktgqos* to a range of IP addresses.

The steps to configure this example are as follows:

1. Configure the switch for the egress mode, by typing the following:
   config qosmode egress
2. Reboot the switch.
3. Create and configure the QoS profile *mktgqos*, by typing the following:
   create qosprofile mktgqos
   config qosprofile mktgqos minbw 0% maxbw 10% priority low
4. Apply the QoS profile to a range of IP addresses, by typing the following:
   config ipqos add 128.1.0.0/16 qosprofile mktgqos

DISPLAYING QOS INFORMATION

To display QoS information on the switch, use the following command:

```
show qosprofile {<qosname> | all}
```

Information displayed includes:

- QoS profile name
- Minimum bandwidth
- Maximum bandwidth
- Priority
- A list of all traffic groups assigned to the QoS profile

Additionally, QoS information can be displayed from the traffic group perspective by using one of the following commands:

- `show fdb permanent`

  Shows destination MAC entries and their QoS profiles.
- `show switch`

  Includes PACE enable/disable information.
- `show vlan`

  Shows the QoS profile assignments to the VLAN.
- `show ipqos`

  Displays the IP QoS table.

RESETTING QOS

To delete a QoS profile use the following command:

```
delete qosprofile <qosname>
```

This command is available only in egress mode.

8 IP Unicast Routing

This chapter describes how to configure IP routing on the BlackDiamond. It assumes that you are already familiar with IP unicast routing. If not, refer to the following publications for additional information:

RFC 1256 — *ICMP Router Discovery Messages*
RFC 1812 — *Requirements for IP Version 4 Routers*

For more information on routing protocols, refer to Chapter 9.

Overview of IP Unicast Routing

The BlackDiamond provides full Layer 3, IP unicast routing. It exchanges routing information with other routers on the network using either the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) protocol. The BlackDiamond dynamically builds and maintains a routing table, and determines the best path for each of its routes.

Each host using the IP unicast routing functionality of the BlackDiamond must have a unique IP address assigned. In addition, the default gateway assigned to the host must be the IP address of the BlackDiamond router interface.

RIP and OSPF are described in Chapter 9.

ROUTER INTERFACES

The routing software and hardware routes IP traffic between router interfaces. A router interface is simply a VLAN that has an IP address assigned to it.

As you create VLANs with IP addresses belonging to different IP subnets, you can also choose to route between the VLANs. Both the VLAN switching and IP routing function occur within the BlackDiamond.

Each IP address and mask assigned to a VLAN must represent a unique IP subnet. You cannot configure the same IP subnet on different VLANs.

In Figure 8-1, a BlackDiamond is depicted with two VLANs defined; *Finance* and *Personnel*. All ports on slots 1 and 3 are assigned to *Finance*; all ports on slots 2 and 4 are assigned to *Personnel*. *Finance* belongs to the IP network 192.207.35.0; the router interface for *Finance* is assigned the IP address 192.206.35.1. *Personnel* belongs to the IP network 192.207.36.0; its router interface is assigned IP address 192.207.36.1. Traffic within each VLAN is switched using the Ethernet MAC addresses. Traffic between the two VLANs is routed using the IP addresses.

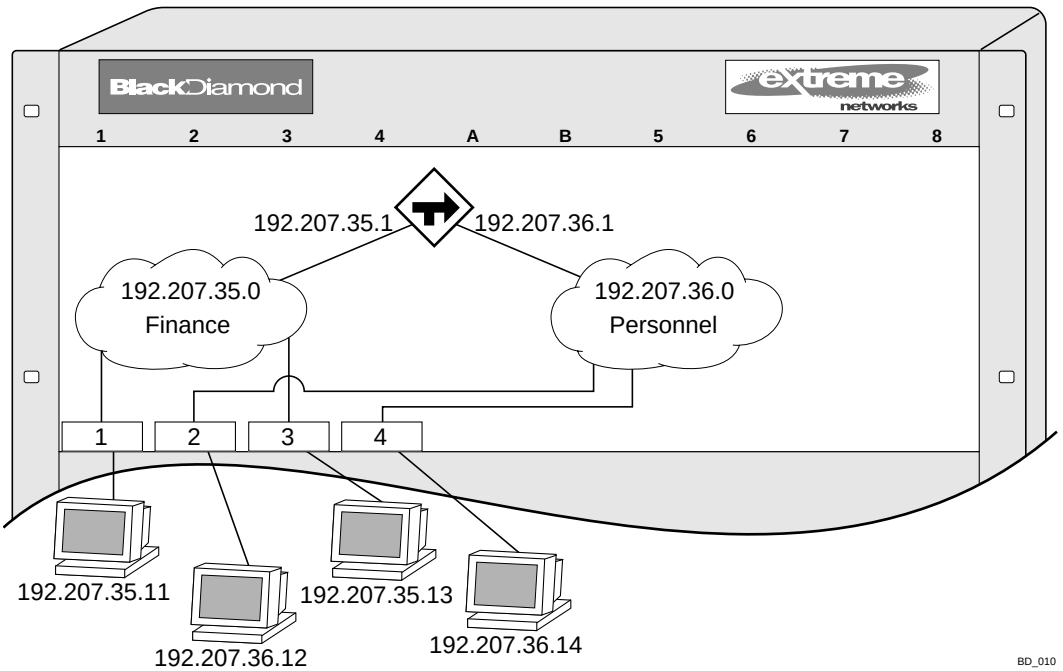

Figure 8-1: Routing between VLANs

POPULATING THE ROUTING TABLE

The BlackDiamond maintains an IP routing table for both network routes and host routes. The table is populated from the following sources:

- Dynamically, by way of routing protocol packets or ICMP redirects exchanged with other routers
- Statically, by way of routes entered by the administrator
  - Default routes, configured by the administrator
  - Locally, by way of interface addresses assigned to the system
  - By other static routes, as configured by the administrator

If you define a default route, and subsequently delete the VLAN on the subnet associated with the default route, the invalid default route entry remains. You must manually delete the configured default route.

DYNAMIC ROUTES

Dynamic routes are typically learned by way of RIP or OSPF. Routers that use RIP or OSPF exchange information in their routing tables in the form of advertisements. Using dynamic routes, the routing table contains only networks that are reachable.

Dynamic routes are aged out of the table when an update for the network is not received for a period of time, as determined by the routing protocol.

STATIC ROUTES

Static routes are manually entered into the routing table. Static routes are used to reach networks not advertised by routers. You can configure up to 64 static unicast routes on the BlackDiamond.

Static routes can also be used for security reasons, to control which routes you want advertised by the router. You can decide if you want all static routes to be advertised, using one of the following commands:

```
[enable | disable] rip export static
[enable | disable] ospf export static
```

The default setting is enabled. Static routes are never aged out of the routing table.

A static route must be associated with a valid IP subnet. An IP subnet is associated with a single VLAN by its IP address and subnet mask. If the VLAN is subsequently deleted, the static route entries using that subnet must be deleted manually.

MULTIPLE ROUTES

When there are multiple, conflicting choices of a route to a particular destination, the router picks the route with the longest matching network mask. If these are still equal, the router picks the route using the following criteria (in the order specified):

- Directly attached network interfaces
- ICMP redirects (refer to Table 8-4)
- Static routes
- Directly attached network interfaces that are not active.

If you define multiple default routes, the route that has the lowest metric is used. If there are multiple default routes that have the same lowest metric, the system picks one of the routes.

You can also configure *blackhole* routes—traffic to these destinations is silently dropped.

PROXY ARP

Proxy ARP was first invented so that ARP-capable devices could respond to ARP Request packets on behalf of ARP-incapable devices. The usage and scope of proxy ARP has been expended since its introduction. Currently, proxy ARP can be used to achieve router redundancy and simplify IP client configuration. The BlackDiamond supports proxy ARP for this type of network configuration. Up to 64 proxy ARP entries can be configured. The section describes some example of how to use proxy ARP with the BlackDiamond.

ARP-INCAPABLE DEVICES

To configure the BlackDiamond to respond to ARP Requests on behalf of devices that are incapable of doing so, you must configure the IP address and MAC address of the ARP-incapable device using the use the following command:

```
config iparp add proxy <ipaddress> {<mask>} <mac_address> {always}
```

Once configured, the system responds to ARP Requests on behalf of the device as long as the following conditions are satisfied:

- The valid IP ARP Request is received on a BlackDiamond router interface.
- The target IP address matches the IP address configured in the proxy ARP table.
- The proxy ARP table entry indicates that the system should always answer this ARP Request (the `always` parameter must be applied).

Once all the proxy ARP conditions are met, the BlackDiamond formulates an ARP Response using the configured MAC address in the packet.

PROXY ARP BETWEEN SUBNETS

In some networks, it is desirable to configure the IP host with a wider subnet than the actual subnet mask of the segment. Proxy ARP can be used so that the router answers ARP Requests for devices outside of the subnet. As a result, the host communicates as if all devices are local. In reality, communication with devices outside of the subnet are proxied by the router.

For example, an IP host is configured with a class B address of 100.101.102.103 and a mask of 255.255.0.0. The BlackDiamond is configured with the IP address 100.101.102.1 and a mask of 255.255.255.0. The BlackDiamond is also configured with a proxy ARP entry of IP address 100.101.0.0 and mask 255.255.0.0, *without* the `always` parameter.

When the IP host tries to communicate with the host at address 100.101.45.67, the IP hosts communicates as if the two hosts are on the same subnet, and sends out an IP ARP Request. The BlackDiamond answers on behalf of the device at address 100.101.45.67, using its own MAC address. All subsequent data packets from 100.101.102.103 are sent to the switch, and the switch routes the packets to 100.101.45.67.

IP MULTINETTING

IP multinetting is used in many legacy IP networks when there is a need to overlap multiple subnets into one physical segment. On the BlackDiamond, you can only assign a single IP address to a router interface (one IP address per VLAN). To support IP multinetting, you must assign multiple VLANs to the same physical port. The BlackDiamond routes IP traffic from one subnet to another, all within the same physical port.

The following rules apply when you are configuring IP multinetting:

- A maximum of one IP address is associated with a router interface (or VLAN).
- Multiple VLANs must be used to implement IP multinetting.
- A maximum of four subnets are allowed on one multinetted port.
- For multinetted segments that span multiple ports, you must configure all the multinetted VLANs with the same port assignment.
- A maximum of one VLAN can run RIP or OSPF, and this VLAN must be configured to use the IP protocol.

BOOTP works only on the VLAN assigned to the IP protocol.

IP MULTINETTING OPERATION

To use IP multinetting, follow these steps:

1 Select a slot and port on which IP multinetting is to run.

For example, slot 1, port 2.

2 Remove the default VLAN from the selected port.

```
config default delete port 1:2
```

3 Create a dummy protocol.

```
create protocol mnet
```

4 Create the multinetted subnets.

```
create vlan net21
create vlan net22
```

5 Assign IP addresses to the net VLANs.

```
config net21 ipaddress 123.45.21.1 255.255.255.0
config net22 ipaddress 192.24.22.1 255.255.255.0
```

6 Assign one of the subnets to the IP protocol.

```
config net21 protocol ip
```

7 Assign the other subnets to the dummy protocol.

```
config net22 protocol mnet
```

8 Assign the subnet to a physical port.

```
config net21 add port 1:2
config net22 add port 1:2
```

9 Enable IP forwarding on the subnets.

```
enable ipforwarding
```

10 Enable IP multinetting.

```
enable multinetting
```

11 If you are using RIP, disable RIP on the dummy VLANs.

```
config rip delete net22
```

IP MULTINETTING EXAMPLES

The following example configures the switch to have one multinetted segment (slot 5, port 5) that contains three subnets (192.67.34.0, 192.67.35.0, and 192.67.37.0).

```
config default delete port 5:5
create protocol mnet
create vlan net34
create vlan net35
create vlan net37
config net34 ipaddress 192.67.34.1
config net35 ipaddress 192.67.35.1
config net37 ipaddress 192.67.37.1
config net34 protocol ip
config net35 protocol mnet
config net37 protocol mnet
config net34 add port 5:5
config net35 add port 5:5
config net37 add port 5:5
enable ipforwarding
enable multinetting
```

The following example configures the switch to have one multinetted segment (slot 5: port 5) that contains three subnets (192.67.34.0, 192.67.35.0, and 192.67.37.0). It also configures a second multinetted segment consisting of two subnets (192.67.36.0 and 192.99.45.0). The second multinetted segment spans three ports (slot1:port 8, slot2:port 9, and slot3:port 10). RIP is enabled on both multinetted segments.

```
config default delete port 5:5
create protocol mnet
create vlan net34
create vlan net35
create vlan net37
config net34 ipaddress 192.67.34.1
config net35 ipaddress 192.67.35.1
config net37 ipaddress 192.67.37.1
config net34 protocol ip
config net35 protocol mnet
config net37 protocol mnet
config net34 add port 5:5
config net35 add port 5:5
config net37 add port 5:5
config default delete port 1:8, 2:9, 3:10
create vlan net36
create vlan net45
config net36 ipaddress 192.67.36.1
config net45 ipaddress 192.99.45.1
config net36 protocol ip
config net45 protocol mnet
config net36 add port 1:8, 2:9, 3:10
config net45 add port 1:8, 2:9, 3:10
config rip delete vlan all
config rip add net34
config rip add net36
enable rip
enable ipforwarding
enable multinetting
```

CONFIGURING IP UNICAST ROUTING

This section describes the commands associated with configuring IP unicast routing on the BlackDiamond. Configuring routing involves the following steps:

1. Create and configure two or more VLANs.

   Although it is possible to enable IP forwarding and an IP routing protocol (such as RIP) with only one VLAN defined, the BlackDiamond does not create or respond appropriately to ICMP messages unless at least two VLANs are created and configured.

For information on creating and configuring VLANs, refer to Chapter 4.

2. Assign each VLAN that will be using routing an IP address, using the following command:

   ```
   config vlan <name> ipaddress <ipaddress> {<mask>}
   ```

   Ensure that each VLAN has a unique IP address.

3. Configure a default route, using the following command:

   ```
   config iproute add default <gateway> {<metric>}
   ```

   Default routes are used when the router has no other dynamic or static route to the requested destination.

4. Turn on IP routing for one or more VLANs, using the following command:

   ```
   enable ipforwarding {vlan <name> | all}
   ```

5. Turn on RIP or OSPF using one of the following commands:

   ```
   enable rip
   enable ospf
   ```

VERIFYING THE IP UNICAST ROUTING CONFIGURATION

Use the `show iproute` command to display the current configuration of IP unicast routing for the system, and for each VLAN. The `show iproute` command displays the currently configured routes, and includes how each route was learned.

Additional verification commands include the following:

- `show iparp`

  Displays the IP ARP table of the system.

- `show ipfdb`

  Displays the hosts that have been transmitting or receiving packets, and the port and VLAN for each host.

- `show ipconfig`

  Displays configuration information for one or more VLANs.

CONFIGURING DHCP/BOOTP RELAY

Once IP unicast routing is configured, you can configure the BlackDiamond to forward Dynamic Host Configuration Protocol (DHCP) or BOOTP requests coming from clients on subnets being service by the BlackDiamond and going to hosts on different subnets. This feature can be used in various applications, including DHCP services between Windows NT servers and clients running Windows 95. To configure the relay function, do the following:

1. Configure VLANs and IP unicast routing.
2. Enable the DHCP or BOOTP relay function, using the following command:

   ```
   enable bootprelay
   ```

3. Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command:

   ```
   config bootprelay add <ipaddress>
   ```

To delete an entry, use the following command:

```
config bootprelay delete {<ipaddress> | all}
```

VERIFYING THE DHCP/BOOTP RELAY CONFIGURATION

To verify the DHCP/BOOTP relay configuration, use the following command:

```
show ipconfig
```

This command displays the configuration of the BOOTP relay service, and the addresses that are currently configured.

USING EXTREME STANDBY ROUTER PROTOCOL

The Extreme Standby Router Protocol (ESRP) provides a way to keep hosts communicating when a physical router becomes unavailable. ESRP works by allowing you to configure a *virtual router* on two or more BlackDiamond systems (and Summit switches) that are part of the same network segment. Connected hosts are then configured using the IP address of the virtual router as the default gateway.

Only one of the virtual routers is active at any given time. This device is called the *active router*. The other BlackDiamond systems and Summit switches remain in standby mode. In the event that the active router fails (determined by the virtual router hello timer), one of the standby routers becomes the active router.

The active router is elected using a priority-based election process. You can configure the priority for each configured device. If the virtual router priority is not configured, the active virtual router is selected by way of the internal MAC address.

You can set up multiple ESRP groups that service different Ethernet segments. You can have a maximum of 32 ESRP groups on your network.

CONFIGURING ESRP

To configure ESRP on the BlackDiamond, create a VLAN that will be used specifically for ESRP.

Table 8-1 describes the commands used to configure ESRP.

Table 8-1: ESRP Commands

Command	Description
enable esrp vlan <name>	Enables ESRP on a VLAN.
config vlan <name> esrp group <number>	Assigns a VLAN to an ESRP group. The valid group numbers are from 1 to 32. The default setting is 1.
config vlan <name> esrp priority <value>	Configures the virtual router priority. The range is 0 to 255. The default setting is 0.

Table 8-1: ESRP Commands (continued)

Command	Description
config vlan <name> esrp timer <hello_timer>	Configures the virtual router hello timer. The range is 1 to 255 seconds. A longer hello timer interval causes the routers to take longer to converge.
disable esrp vlan <name>	Disables ESRP on a VLAN.

The following example creates a VLAN named *esrp1*.

```
create vlan esrp1
config vlan esrp1 add port 2:1-2:4
config vlan esrp1 ipaddress 10.1.2.3/16
config vlan esrp1 timer 10
config vlan esrp1 esrp priority 10
enable esrp vlan esrp1
```

IP COMMANDS

Table 8-2 describes the commands used to configure basic IP settings.

Table 8-2: Basic IP Commands

Command	Description
enable bootp vlan [<name> | all]	Enables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server. The default setting is enabled for all VLANs.
enable bootprelay	Enables the forwarding of BOOTP and Dynamic Host Configuration Protocol (DHCP) requests.
enable ipforwarding {vlan <name> | all}	Enables IP routing for one or more VLANs. If no argument is provided, enables routing for all VLANs that have been configured with an IP address. The default setting for `ipforwarding` is disabled.

Table 8-2: Basic IP Commands (continued)

Command	Description
enable ipforwarding broadcast {vlan <name> | all}	Enables forwarding IP broadcast traffic for one or more VLANs. If no argument is provided, enables broadcast forwarding for all VLANs. To enable, `ipforwarding` must be enabled on the VLAN. The default setting is enabled.
enable multinetting	Enables IP multinetting on the system.
config bootprelay add <ipaddress>	Adds the IP destination address to forward BOOTP packets.
config bootprelay delete [<ipaddress> | all]	Removes one or all IP destination addresses for forwarding BOOTP packets.
config iparp add <ipaddress> <mac_address>	Adds a permanent entry to the ARP table. Specify the IP address and MAC address of the entry.
config iparp delete <ipaddress>	Deletes an entry from the ARP table. Specify the IP address of the entry.
disable bootp vlan [<name> | all]	Disables the generation and processing of BOOTP packets.
config iparp add proxy <ipaddress> {<mask>} {<mac_address>} {always}	Configures proxy ARP entries. Up to 64 proxy ARP entries can be configured. When `mask` is not specified, an address with the mask 255.255.255.255 is assumed. When `mac_address` is not specified, the MAC address of the switch is used in the ARP Response. When `always` is specified, the switch answers ARP Requests without filtering requests that belong to the same subnet of the receiving router interface.
config iparp delete proxy [<ipaddress> {<mask>} | all]	Deletes one or all proxy ARP entries.
disable bootprelay	Disables the forwarding of BOOTP requests.
disable ipforwarding {vlan <name> | all}	Disables routing for one or more VLANs.
disable ipforwarding broadcast {vlan <name> | all}	Disables routing of broadcasts to other networks.
disable multinetting	Disables IP multinetting on the system.
clear iparp [<ipaddress> <mask> | vlan <name> | all]	Removes dynamic entries in the IP ARP table. Permanent IP ARP entries are not affected.

Table 8-2: Basic IP Commands (continued)

Command	Description
clear ipfdb [<ipaddress> | vlan <name> | all]	Removes the dynamic entries in the IP forwarding database.
enable esrp port <slot:portlist>	Enables ESRP on a port.
config port <slot:portlist> esrp group <number>	Assigns a port to an ESRP group. The valid group numbers are from 1 to 32.
config port <slot:portlist> esrp priority <value>	Configures the virtual router priority. The range is 1to 32.
config port <slot:portlist> esrp timer <hello_value>	Configures the virtual router keep-alive timer. The range is 1 to 255.
disable esrp port <slot:portlist>	Disables ESRP on a port.

Table 8-3 describes the commands used to configure the IP route table.

Table 8-3: Route Table Configuration Commands

Command	Description
enable iproute sharing	Enables load sharing if multiple routes to the same destination are available. Only paths with the same lowest cost are shared. The default setting is enabled.
config ipqos add <ip_destination_address> <mask> qosprofile <qosname>	Adds a QoS profile to an IP destination address.
config ipqos delete <ip_destination_address> <mask>	Deletes a QoS profile from an IP destination address.
config iproute add <ipaddress> <mask> <gateway> <metric>	Adds a static address to the routing table. Use a value of 255.255.255.255 for `mask` to indicate a host entry
config iproute delete <ipaddress> <mask> <gateway>	Deletes a static address from the routing table.
config iproute add blackhole <ipaddress> <mask>	Adds a `blackhole` address to the routing table. All traffic destined for the configured IP address is dropped, and no Internet Control Message Protocol (ICMP) message is generated.
config iproute delete blackhole <ipaddress> <mask>	Deletes a `blackhole` address from the routing table.

Table 8-3: Route Table Configuration Commands (continued)

Command	Description
config iproute add default <gateway> {<metric>}	Adds a default gateway to the routing table. A default gateway must be located on a configured IP interface. If no metric is specified, the default metric of 1 is used.
config iproute delete default <gateway>	Deletes a default gateway from the routing table.
disable iproute sharing	Disables load sharing for multiple routes.

Table 8-4 describes the commands used to configure the ICMP protocol.

Table 8-4: ICMP Configuration Commands

Command	Description
enable icmp redirects {vlan <name> | all}	Enables generation of ICMP redirect messages on one or more VLANs. The default setting is enabled.
enable icmp unreachables {vlan <name> | all}	Enables the generation of ICMP unreachable messages on one or more VLANs. The default setting is enabled.
enable icmp useredirects	Enables the modification of route table information when an ICMP redirect message is received. The default setting is disabled.
enable irdp {vlan <name> | all}	Enables the generation of ICMP router advertisement messages on one or more VLANs. The default setting is enabled.
config irdp [multicast | broadcast]	Configures the destination address of the router advertisement messages. The default setting is multicast.

Table 8-4: ICMP Configuration Commands (continued)

Command	Description
config irdp <mininterval> <maxinterval> <lifetime> <preference>	Configures the router advertisement message timers, using seconds. Specify: ■ `mininterval` — The minimum amount of time between router advertisements. The default setting is 450 seconds. ■ `maxinterval` — The maximum time between router advertisements. The default setting is 600 seconds. ■ `lifetime` — The default setting is 1,800 seconds. ■ `preference` — The preference level of the router. An IRDP client always uses the router with the highest preference level. Change this setting to encourage or discourage the use of this router. The default setting is 0.
unconfig icmp	Resets all ICMP settings to the default values.
unconfig irdp	Resets all router advertisement settings to the default values.
disable icmp redirects {vlan <name> | all}	Disables the generation of ICMP redirects on one or more VLANs.
disable icmp unreachables {vlan <name> | all}	Disables the generation of ICMP unreachable messages on one or more VLANs.
disable icmp useredirects	Disables the changing of routing table information when an ICMP redirect message is received.
disable irdp {vlan <name> | all}	Disables the generation of router advertisement messages on one or more VLANs.

ROUTING CONFIGURATION EXAMPLE

Figure 8-2 illustrates a switch that has three VLANs defined as follows:

- *Finance*
  - Protocol-sensitive VLAN using the IP protocol
  - All ports on slots 1 and 3 have been assigned
  - IP address 192.207.35.1

- *Personnel*
  - — Protocol-sensitive VLAN using the IP protocol
  - — All ports on slots 2 and 4 have been assigned
  - — IP address 192.207.36.1
- *MyCompany*
  - — Port-based VLAN
  - — All ports on slots 1 through 4 have been assigned

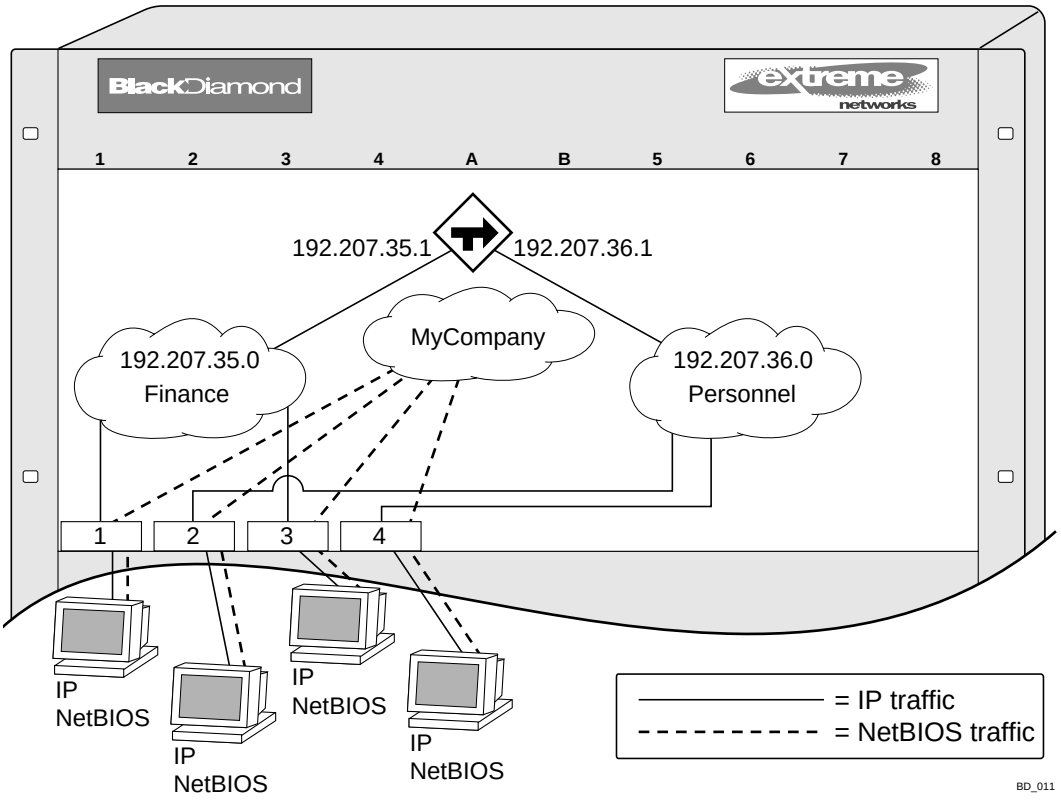

Figure 8-2: Unicast routing configuration example

The stations connected to the system generate a combination of IP traffic and NetBIOS traffic. The IP traffic is filtered by the protocol-sensitive VLANs. All other traffic is directed to the VLAN *MyCompany*.

In this configuration, all IP traffic from stations connected to slots 1 and 3 have access to the router by way of the VLAN *Finance*. Ports on slots 2 and 4 reach the router by way of the VLAN *Personnel*. All other traffic (NetBIOS) is part of the VLAN *MyCompany*.

The example in Figure 8-2 is configured as follows:

```
create vlan Finance
create vlan Personnel
create vlan MyCompany

config Finance protocol ip
config Personnel protocol ip

config Finance add port 1:*,3:*
config Personnel add port 2:*,4:*
config MyCompany add port all

config Finance ipaddress 192.207.35.1
config Personnel ipaddress 192.207.36.1

config rip add vlan Finance
config rip add vlan Personnel

enable ipforwarding
enable rip
```

DISPLAYING ROUTER SETTINGS

To display settings for various IP routing components, use the commands listed in Table 8-5.

Table 8-5: Router Show Commands

Command	Description
show iparp proxy {<ipaddress> {<mask>} | all}	Displays the proxy ARP table.

Table 8-5: Router Show Commands (continued)

Command	Description
show ipconfig {vlan <name> | all}	Displays configuration information for one or more VLANs, including the following: ■ IP address, subnet mask ■ IP forwarding information ■ BOOTP configuration ■ VLAN name, VLANid ■ Global ICMP configuration ■ Global IGMP configuration ■ Global router advertisement configuration
show ipqos {<ip_destination_address> <mask> | all}	Displays the IP QoS table.
show ipstats {vlan [<name> | all]}	Displays IP statistics for the CPU of the system.
show iparp {<ipaddress | vlan <name> | all | permanent}	Displays the IP Address Resolution Protocol (ARP) table. You can filter the display by IP address, VLAN, or permanent entries. Each entry displayed includes the following: ■ IP address ■ MAC address ■ Aging timer value ■ VLAN name, VLANid, and port number ■ Flags
show ipfdb {<ipaddress> <netmask> | vlan <name> | all}	Displays the contents of the IP forwarding database table. Used for technical support purposes.
show iproute vlan {<name> | all | permanent | <ipaddress> <mask>}	Displays the contents of the IP routing table.

RESETTING AND DISABLING ROUTER SETTINGS

To return router settings to their defaults and disable routing functions, use the commands listed in Table 8-6.

Table 8-6: Router Reset and Disable Commands

Command	Description
clear iparp [<ipaddress> | vlan <name> | all]	Removes dynamic entries in the IP ARP table. Permanent IP ARP entries are not affected.
clear ipfdb [<ipaddress> <netmask> | vlan <name> | all]	Removes the dynamic entries in the IP forwarding database.
disable bootp vlan [<name> | all]	Disables the generation and processing of BOOTP packets.
disable bootprelay	Disables the forwarding of BOOTP requests.
disable icmp redirects {vlan <name> | all}	Disables the generation of ICMP redirects on one or more VLANs.
disable icmp unreachables	Disables the generation of ICMP unreachable messages on one or more VLANs.
disable icmp useredirects	Disables the changing of routing table information when an ICMP redirect message is received.
disable ipforwarding {vlan <name> | all}	Disables routing for one or more VLANs.
disable ipforwarding broadcast {vlan <name> | all}	Disables routing of broadcasts to other networks.
disable irdp {vlan <name> | all}	Disables the generation of router advertisement messages on one or more VLANs.
unconfig icmp	Resets all ICMP settings to the default values.
unconfig irdp	Resets all router advertisement settings to the default values.

9 Routing Protocols

This chapter describes the IP unicast routing protocols available on the BlackDiamond. It assumes that you are already familiar with IP unicast routing. If not, refer to the following publications for additional information:

RFC 1058 — *Routing Information Protocol (RIP)*
RFC 1256 — *ICMP Router Discovery Messages*
RFC 1723 — *RIP Version 2*
RFC 2178 — *OSPF Version 2*

Interconnections: Bridges and Routers
by Radia Perlman
ISBN 0-201-56332-0
Published by Addison-Wesley Publishing Company

OVERVIEW

The BlackDiamond supports the use of the Routing Information Protocol (RIP) and the Open Shortest Path First (OSPF) protocol for IP unicast routing.

RIP is a distance-vector protocol, based on the Bellman-Ford (or distance-vector) algorithm. The distance-vector algorithm has been in use for many years, and is widely deployed and understood.

OSPF is a link-state protocol, based on the Dijkstra link-state algorithm. OSPF is a newer Interior Gateway Protocol (IGP), and solved a number of problems associated with using RIP on today's complex networks.

RIP VERSUS OSPF

The distinction between RIP and OSPF lies in the fundamental differences between distance-vector protocols and link-state protocols. Using a distance-vector protocol, each router creates a unique routing table from summarized information obtained from neighboring routers. Using a link-state protocol, every router maintains an identical routing table created from information obtained from all routers in the autonomous system. Each router builds a shortest path tree, using itself as the root. The link-state protocol ensures that updates sent to neighboring routers are acknowledged by the neighbors, verifying that all routers have a consistent network map.

The biggest advantage of using RIP is that it is relatively simple to understand and implement, and it has been the *de facto* routing standard for many years.

RIP has a number of limitations that can cause problems in large networks, including the following:

- A limit of 15 hops between the source and destination networks
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table
- Slow convergence
- Routing decisions based on hop count; no concept of link costs or delay
- Flat networks; no concept of areas or boundaries

OSPF offers many advantages over RIP, including the following:

- No limitation on hop count
- Route updates multicast only when changes occur
- Faster convergence
- Support for load balancing to multiple routers based on the actual cost of the link
- Support for hierarchical topologies where the network is divided into areas

The details of RIP and OSPF are explained later in this chapter.

OVERVIEW OF RIP

RIP is an Interior Gateway Protocol (IGP) first used in computer routing in the ARPAnet as early as 1969. It is primarily intended for use in homogeneous networks of moderate size.

To determine the best path to a distant network, a router using RIP always selects the path that has the least number of hops. Each router that data must traverse is considered to be one hop.

ROUTING TABLE

The routing table in a router using RIP contains an entry for every known destination network. Each routing table entry contains the following information:

- IP address of the destination network
- Metric (hop count) to the destination network
- IP address of the next router
- Timer that tracks the amount of time since the entry was last updated

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

SPLIT HORIZON

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

POISON REVERSE

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same port that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

TRIGGERED UPDATES

Triggered updates occur whenever a router changes the metric for a route, and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIP-related traffic.

ROUTE ADVERTISEMENT OF VLANS

VLANs that are configured with an IP address, but are configured to not route IP, have their subnets advertised by RIP with a metric of 16 (unreachable). To disable the advertising of a subnet completely, you must unconfigure the IP address for the VLAN using the following command:

```
unconfig vlan <name> ipaddress
```

RIP VERSION 1 VERSUS RIP VERSION 2

A new version of RIP, called RIP version 2, expands the functionality of RIP version 1 to include the following:

- Variable-Length Subnet Masks (VLSMs)
- Next-hop addresses

Support for next-hop addresses allows for optimization of routes in certain environments.

- Multicasting

RIP version 2 packets can be multicast instead of being broadcast, reducing the load on hosts that do not support routing protocols.

OVERVIEW OF OSPF

OSPF is a link-state protocol that distributes routing information between routers belonging to a single IP domain, also known as an *autonomous system*. In a link-state routing protocol, each router maintains a database describing the topology of the autonomous system. Each participating router has an identical database maintained from the perspective of that router.

From the link-state database (LSDB), each router constructions a tree of shortest paths, using itself as the root. The shortest path tree provides the route to each destination in the autonomous system. When several equal-cost routes to a destination exist, traffic can distributed among them. The cost of a route is described by a single metric.

LINK-STATE DATABASE

Upon initialization, each router transmits a link-state advertisement (LSA) on each of its interfaces. The LSA contains the following information for each link:

- IP network number of the link
- Subnet mask of the link
- Metric for the link
- Operation status (up or down) of the link

LSAs are collected by each router and entered into the LSDB of each router. OSPF uses flooding to distribute LSAs between routers. Any change in routing information is sent to all of the routers in the network. All routers within an area have the exact same LSDB.

AREAS

OSPF allows parts of a networks to be grouped together into *areas*. The topology within an area is hidden from the rest of the autonomous system. Hiding this information enables a significant reduction in LSA traffic, and reduces the computations needed to maintain the LSDB. Routing within the area is determined only by the topology of the area.

The three types of routers defined by OSPF are as follows:

- Internal Router (IR)

  An internal router has all of its interfaces within the same area.

- Area Border Router (ABR)

  An ABR has interfaces in multiple areas. It is responsible for exchanging summary advertisements with other ABRs.

- Autonomous System Border Router (ASBR)

  An ASBR acts as a gateway between OSPF and other routing protocols, or other autonomous systems.

AREA 0

Any OSPF network that contains more than one area is required to have an area configured as area 0, also called the *backbone*. All areas in an autonomous system must be connected to the backbone. When designing networks, you should start with area 0, and then expand into other areas.

The backbone allows summary information to be exchanged between ABRs. Every ABR hears the area summaries from all other ABRs. The ABR then forms a picture of the distance to all network outside of its area by examining the collected advertisements, and adding in the backbone distance to each advertising router.

STUB AREAS

OSPF allows certain areas to be configured as *stub areas*. A stub area is connected to only one other area and contains a single exit point. The area that connects to a stub area can be the backbone area. All routing out of a stub area is based on default routes. Stub areas are used to reduce memory and computation requirements on OSPF routers. Figure 9-1 shows a stub area.

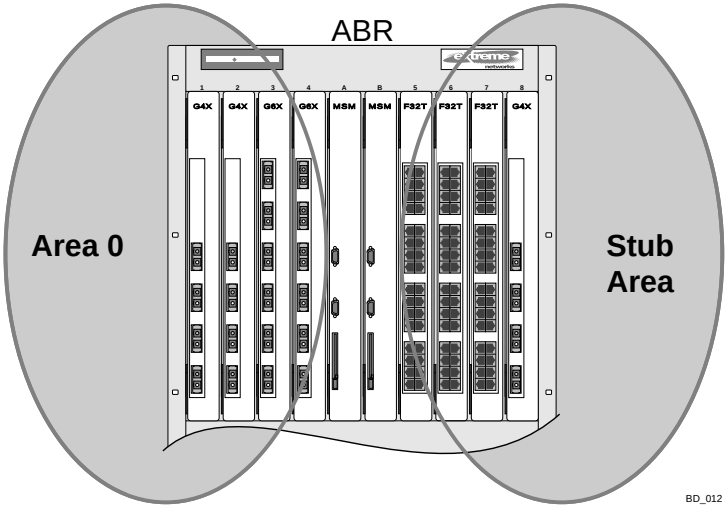

Figure 9-1: Stub area

VIRTUAL LINKS

In the situation when a new area is introduced that does have a direct physical attachment to the backbone, a *virtual link* is used. A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Figure 9-2 illustrates a virtual link.

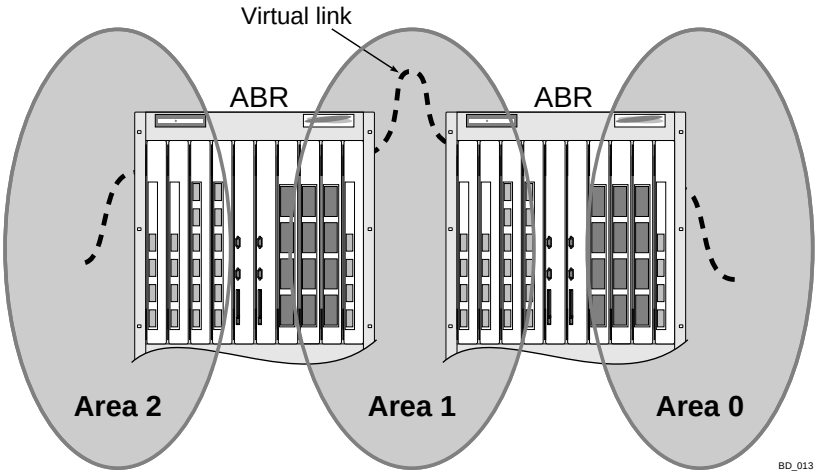

Figure 9-2: Virtual link for stub area

Virtual links are also used to repair a discontiguous backbone area. For example, in Figure 9-3, if the connection between ABR1 and the backbone fails, the connection using ABR2 provides redundancy so that the discontiguous area can continue to communicate with the backbone using the virtual link.

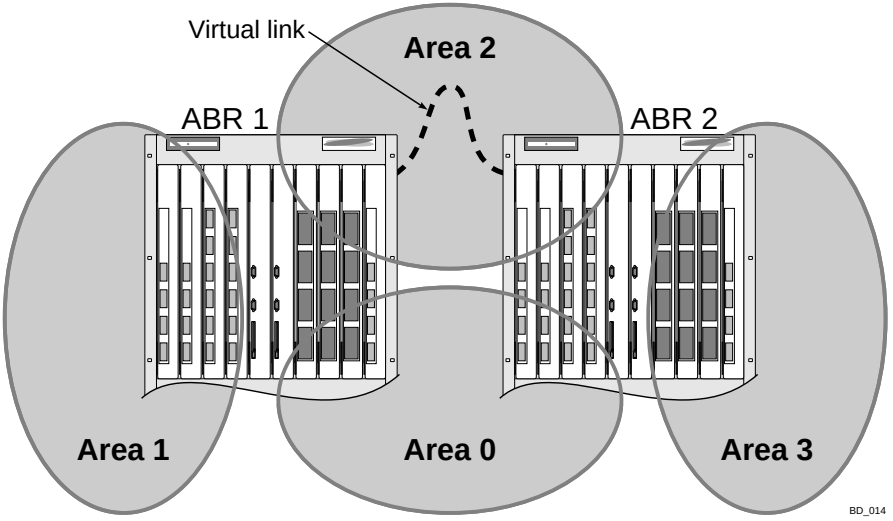

Figure 9-3: Virtual link providing redundancy

UNICAST ROUTE DISTRIBUTION

Both RIP and OSPF can be enabled simultaneously on the BlackDiamond system. If you enable both protocols, you have the option of propagating routes learned by one protocol to the other protocol. This is known as *route distribution*.

Using route distribution, you can configure the following options:

- Enable RIP to export OSPF routes, and configure the metric, using the command

  `enable rip export ospf {<metric>}`
- Enable OSPF to export RIP routes, and configure the OSPF LSA type and metric, using the command

  `enable ospf export rip type [1 | 2] {<metric>}`

CONFIGURING RIP

Table 9-1 describes the commands used to configure RIP.

Table 9-1: RIP Configuration Commands

Command	Description
enable rip	Enables RIP. The default setting is disabled.
enable rip aggregation	Enables RIP aggregation of subnet information an interface configured to sent RIP v2 or RIP v2-compatible traffic. The BlackDiamond summarizes subnet routes to the nearest class network route. The following rules apply when using RIP aggregation: ■ Subnet routes are aggregated to the nearest class network route when crossing a class boundary. ■ Within a class boundary, no routes are aggregated. ■ If aggregation is enabled, the behavior is the same as in RIP v1. ■ If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary. The default setting is enabled.

Table 9-1: RIP Configuration Commands (continued)

Command	Description
enable rip export static	Enables the advertisement of static routes using RIP. The default setting is enabled.
enable rip export ospf {<metric>}	Enables the distribution of OSPF routes into the RIP domain. The default setting is disabled.
enable rip poisonreverse	Enables the split horizon with poison-reverse algorithm for RIP. The default setting is enabled. If you enable poison reverse and split horizon, poison reverse takes precedence.
enable rip splithorizon	Enables the split horizon algorithm for RIP. Default setting is enabled.
enable rip triggerupdate	Enables triggered updates. *Triggered update*s are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes, or changes the metric of a route. The default setting is enabled.
config rip add {vlan <name> | all}	Configures RIP on an IP interface. If no VLAN is specified, then `all` is assumed. When an IP interface is created, per interface RIP configuration is disabled by default.
config rip delete [vlan <name> | all]	Disables RIP on an IP interface. When RIP is disabled on the interface, the parameters are not reset to their defaults.
config rip garbagetime {<delay>}	Configures the RIP garbage time. The timer granularity is 10 seconds. The default setting is 120 seconds.
config rip routetimeout {<delay>}	Configures the route timeout. The timer granularity is 10 seconds. The default setting is 180 seconds.
config rip rxmode [none | v1only | v2only | any] {vlan <name> | all}	Changes the RIP receive mode for one or more VLANs. Specify: ■ `none` — Drop all received RIP packets. ■ `v1only` — Accept only RIP v1 format packets. ■ `v2only` — Accept only RIP v2 format packets. ■ `any` — Accept both RIP v1 and v2 packets. If no VLAN is specified, the setting is applied to all VLANs. The default setting is `any`.

Table 9-1: RIP Configuration Commands (continued)

Command	Description
config rip txmode [none | v1only | v1comp | v2only] {vlan <name> | all}	Changes the RIP transmission mode for one or more VLANs. Specify: ■ `none` — Do not transmit any packets on this interface. ■ `v1only` — Transmit RIP v1 format packets to the broadcast address. ■ `v1comp` — Transmit RIP v2 format packets to the broadcast address. ■ `v2only` — Transmit RIP v2 format packets to the RIP multicast address If no VLAN is specified, the setting is applied to all VLANs. The default setting is `v2only`.
config rip updatetime {<delay>}	Changes the periodic RIP update timer. The timer granularity is 10 seconds. The default setting is 30 seconds.

RIP CONFIGURATION EXAMPLE

Figure 9-4 illustrates a switch that has three VLANs defined as follows:

- *Finance*
  - Protocol-sensitive VLAN using the IP protocol
  - All ports on slots 1 and 3 have been assigned
  - IP address 192.207.35.1
- *Personnel*
  - Protocol-sensitive VLAN using the IP protocol
  - All ports on slots 2 and 4 have been assigned
  - IP address 192.207.36.1
- *MyCompany*
  - Port-based VLAN
  - All ports on slots 1 through 4 have been assigned

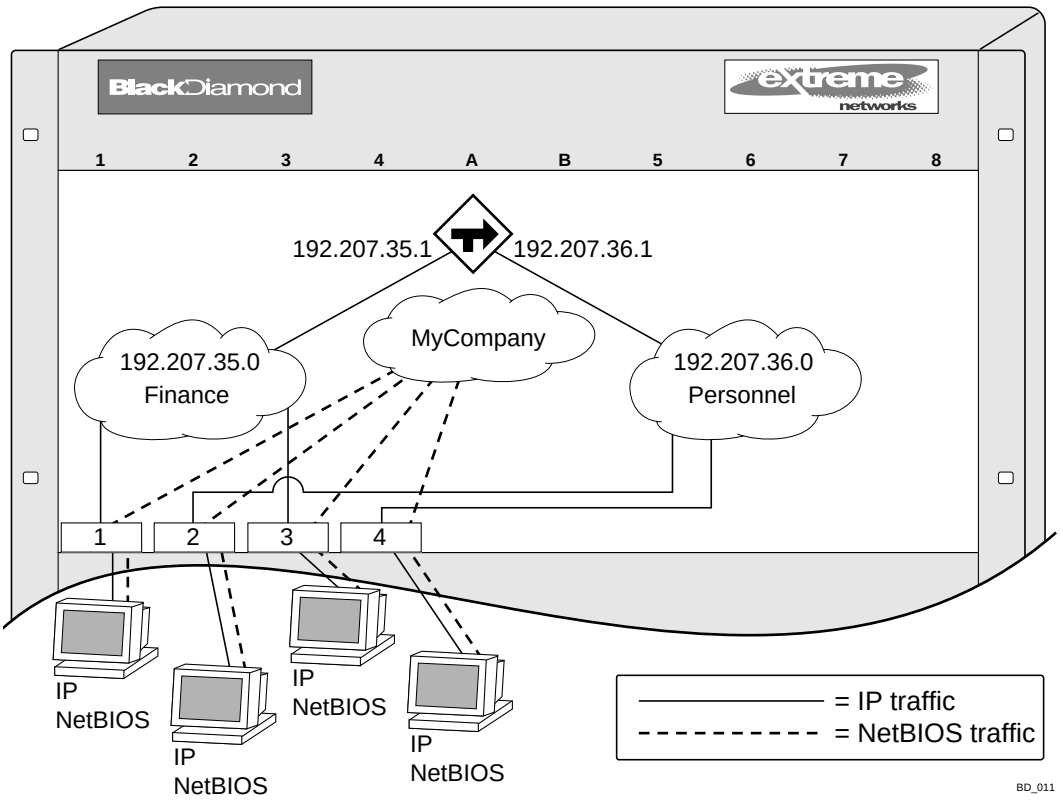

Figure 9-4: RIP configuration example

The stations connected to the system generate a combination of IP traffic and NetBIOS traffic. The IP traffic is filtered by the protocol-sensitive VLANs. All other traffic is directed to the VLAN *MyCompany.*

In this configuration, all IP traffic from stations connected to slots 1 and 3 have access to the router by way of the VLAN *Finance.* Ports on slots 2 and 4 reach the router by way of the VLAN *Personnel*. All other traffic (NetBIOS) is part of the VLAN *MyCompany.*

The example in Figure 9-4 is configured as follows:

```
create vlan Finance
create vlan Personnel
create vlan MyCompany

config Finance protocol ip
config Personnel protocol ip

config Finance add port 1:*,3:*
config Personnel add port 2:*,4:*
config MyCompany add port all

config Finance ipaddress 192.207.35.1
config Personnel ipaddress 192.207.36.1

enable ipforwarding
config rip add vlan all
enable rip
```

DISPLAYING RIP SETTINGS

To display settings for RIP, use the commands listed in Table 9-2.

Table 9-2: RIP Show Commands

Command	Description
show rip {vlan <name> | all}	Displays RIP configuration and statistics for one or more VLANs.
show rip stat {vlan <name> | all}	Displays RIP-specific statistics. Statistics include the following per interface: ■ Packets transmitted ■ Packets received ■ Bad packets received ■ Bad routes received ■ Number of RIP peers ■ Peer information

RESETTING AND DISABLING RIP

To return RIP settings to their defaults, or to disable RIP, use the commands listed in Table 9-3.

Table 9-3: RIP Reset and Disable Commands

Command	Description
config rip delete [vlan <name> | all]	Disables RIP on an IP interface. When RIP is disabled on the interface, the parameters are not reset to their defaults.
disable rip	Disables RIP.
disable rip aggregation	Disables the RIP aggregation of subnet information on a RIP v2 interface.
disable rip splithorizon	Disables split horizon.
disable rip poisonreverse	Disables poison reverse.
disable rip triggerupdate	Disables triggered updates.
disable rip export static	Disables the filtering of static routes.
disable rip export osfp	Disables the distribution of OSPF routes into the RIP domain.
unconfig rip {vlan <name> | all}	Resets all RIP parameters to the default VLAN. Does not change the enable/disable state of the RIP settings.

CONFIGURING OSPF

Table 9-4 describes the commands used to configure OSPF.

Table 9-4: OSPF Configuration Commands

Command	Description
create ospf area <areaid>	Creates an OSPF area. By default, the OSPF area 0.0.0.0 is created.
enable ospf	Enables OSPF process for the router.
enable ospf export static type [1 | 2]	Exports statically configured routes to other OSPF routers. The default setting is enabled.

Table 9-4: OSPF Configuration Commands (continued)

Command	Description
enable ospf export rip type [1 | 2] {<metric>}	Allows OSPF to export RIP routes using LSAs to other OSPF routers. The default setting is disabled.
config ospf [vlan <name> | area <areaid> | virtual-link <routerid> <areaid>] authentication [simple-password <password> | md5 <md5_key_id> <md5_key>| none]	Specifies the authentication password (up to 8 characters) or MD5 key for one or all interfaces in an area. The `md5_key` is a numeric value with the range 0 to 65536. When the OSPF area is specified, authentication information is applied to all OSPF interfaces within the area.
config ospf vlan <name> area <areaid>	Associates a VLAN (router interface) with an OSPF area. All router interfaces must have an associated OSPF area. The default `areaid` is 0 (backbone area).
config ospf [vlan <name> | area <areaid> | all] cost <number>	Configures the cost metric of one or all interface(s). The default cost is 1.
config ospf [vlan <name> | area <areaid> | all] priority <number>	Configures the priority used in the designated router election algorithm for one or all IP interface(s) of for all the interfaces within the area. The range is 0 through 255, and the default setting is 1.
config ospf add [vlan <name> | all]	Enables OSPF on one or all VLANs (router interfaces). The default setting is disabled.
config ospf delete [vlan <name> | all]	Disables OSPF on one or all VLANs (router interfaces).
config ospf add virtual-link <routerid> <areaid>	Adds a virtual link connected to another ABR. Specify the following: ■ `routerid` — Far-end router interface number. ■ `areaid` — Transit area used for connecting the two end-points. The transit area cannot have the IP address 0.0.0.0.
config ospf delete virtual-link <routerid> <areaid>	Removes a virtual link.
config ospf area <areaid> normal	Configures an OSFP area as a normal area. The default setting is `normal`.
config ospf area <areaid> stub [summary | nosummary] stub-default-cost <cost>	Configures an OSPF area as a stub area.
config ospf area add range <ipaddress> <mask> [advertise | noadvertise]	Configures a range of IP addresses in an OSPF area. If advertised, the range is exported as a single summary LSA by the ABR.

Table 9-4: OSPF Configuration Commands (continued)

Command	Description
config ospf area delete range <ipaddress> <mask>	Deletes a range of IP addresses in an OSPF area.
config ospf routerid [automatic | <routerid>]	Configures the OSPF router ID. If automatic is specified, the switch uses the largest IP interface address as the OSPF router ID. The default setting is `automatic`.
config ospf [vlan <name> | area <areaid> | virtual-link <routerid>] timer <retransmission_interval> <transmission_delay> <hello_interval> <dead_interval>	Configures the timers for one interface or all interfaces in the same OSPF area. The following default, minimum, and maximum values (in seconds) are used: ■ `retransmission_interval` Default: 5 Minimum: 0 Maximum: 3600 ■ `transmission_delay` Default: 1 Minimum: 0 Maximum: 3600 ■ `hello _interval` Default: 10 Minimum: 1 Maximum: 65535 ■ `dead_interval` Default: 40 Minimum: 1 Maximum: 2147483647
config ospf spf-hold-time {<seconds>}	Configures the minimum number of seconds between Shortest Path First (SPF) recalculations. The default setting is 3 seconds.

OSPF CONFIGURATION EXAMPLE

Figure 9-5 shows an example of an autonomous system using OSPF routers. The details of this network follow.

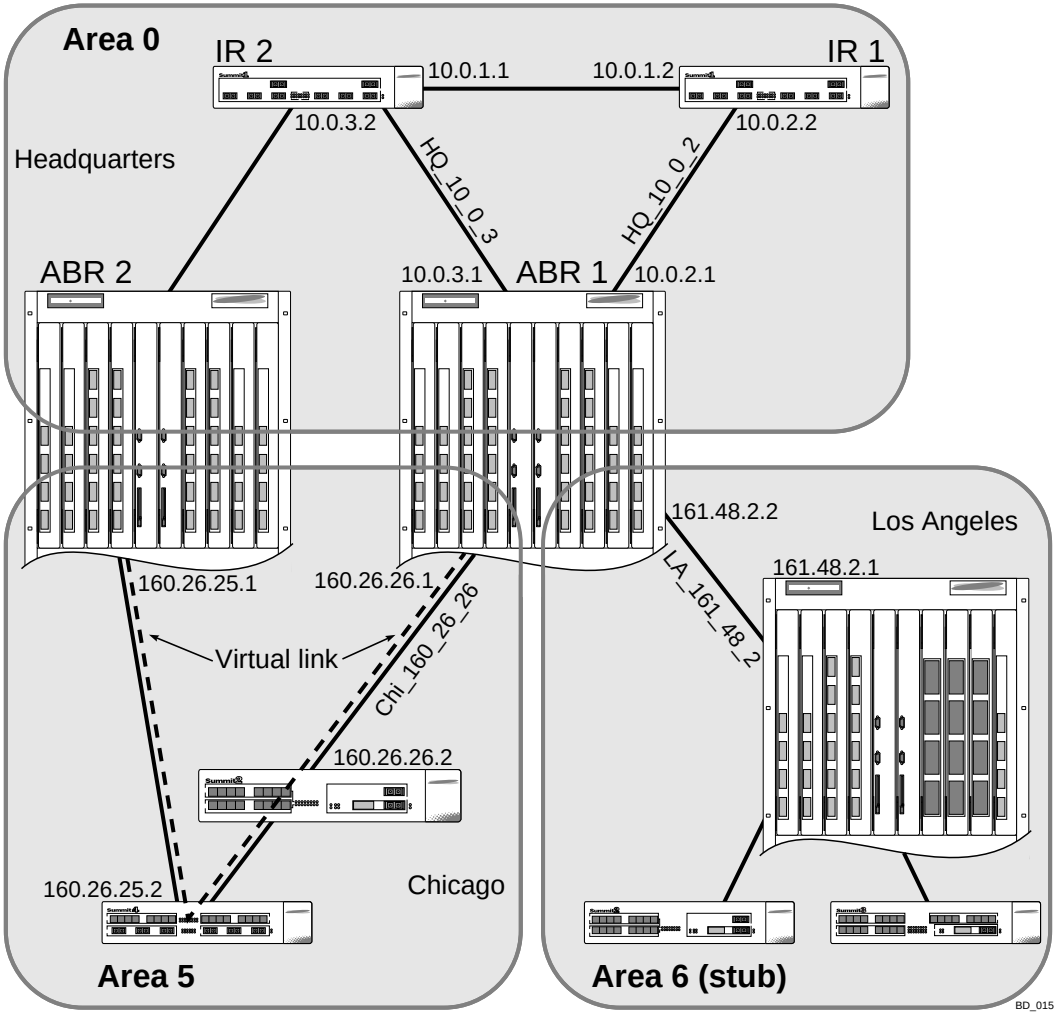

Figure 9-5: OSPF configuration example

Area 0 is the backbone area. It is located at the headquarters and has the following characteristics:

- 2 internal routers (IR1 and IR2)
- 2 area border routers (ABR1 and ABR2)
- Network number 10.0.x.x
- 2 identified VLANs (HQ_10_0_2 and HQ_10_0_3)

Area 5 is connected to the backbone area by way of ABR1 and ABR2. It is located in Chicago and has the following characteristics:

- Network number 160.26.x.x
- 1 identified VLAN (Chi_160_26_26)
- 2 internal routers
- A virtual link from ABR1 to ABR2 that traverses both internal routers.

  In the event that the link between either ABR and the backbone fails, the virtual link provides a connection for all routers that become discontiguous from the backbone.

Area 6 is a stub area connected to the backbone by way of ABR1. It is located in Los Angeles and has the following characteristics:

- Network number 161.48.x.x
- 1 identified VLAN (LA_161_48_2)
- 3 internal routers
- Uses default routes for inter-area routing

Two router configurations for the example in Figure 9-5 are provided in the following section.

CONFIGURATION FOR ABR1

The following is the configuration for the router labeled ABR1:

```
create vlan HQ_10_0_2
create vlan HQ_10_0_3
create vlan LA_161_48_2
create vlan Chi_160_26_2

config vlan HQ_10_0_2 ipaddress 10.0.2.1 255.255.255.0
config vlan HQ_10_0_3 ipaddress 10.0.3.1 255.255.255.0
config vlan LA_161_48_2 ipaddress 161.48.2.2 255.255.255.0
config vlan Chi_160_26_2 ipaddress 160.26.2.1 255.255.255.0

create ospf area 0.0.0.5
create ospf area 0.0.0.6

enable ipforwarding

config ospf area 0.0.0.6 stub nosummary stub-default-cost 10
config ospf vlan LA_161_48_2 area 0.0.0.6
config ospf vlan Chi_160_26_2 area 0.0.0.5
config ospf add virtual-link 160.26.25.1 0.0.0.5
config ospf add vlan all

enable ospf
```

CONFIGURATION FOR IR1

The following is the configuration for the router labeled IR1:

```
config vlan HQ_10_0_1 ipaddress 10.0.1.2 255.255.255.0
config vlan HQ_10_0_2 ipaddress 10.0.2.2 255.255.255.0
config ospf add vlan all
enable ipforwarding
enable ospf
```

DISPLAYING OSPF SETTINGS

To display settings for OSPF, use the commands listed in Table 9-5.

Table 9-5: OSPF Show Commands

Command	Description
show ospf	Displays global OSPF information.
show ospf area {<areaid> | all}	Displays information about a particular OSPF area, or all OSPF areas.
show ospf interfaces {vlan <name> | area <areaid> | all}	Displays information about one or all OSPF interfaces. If no argument is specific, all OSPF interfaces are displayed.
show ospf lsdb {detail} {area <areaid> | all} {router | network | summary_net | summary_asb | as_external | all}	Displays a table of the current LSDB. You can filter the display using either the area ID or the remote router's router ID, or the link-state ID. The default is `all` with no detail. If detail is specified, each entry includes complete LSA information.
show ospf virtual-link {<areaid> <routerid> | all}	Displays virtual link information about a particular router or all routers.

RESETTING AND DISABLING OSPF SETTINGS

To return OSPF settings to their defaults, use the commands listed in Table 9-6.

Table 9-6: OSPF Reset and Disable Commands

Command	Description
config ospf delete [vlan <name> | all]	Disables OSPF on one or all VLANs (router interfaces).
delete ospf area [<areaid> | all]	Deletes and OSPF area. Once an OSPF area is removed, the associated OSPF area and OSPF interface information is removed.
disable ospf	Disables OSPF.
disable ospf export static	Disables exporting of statically configured routes.
disable ospf export rip	Disables exporting of RIP routes to other OSPF routers.

10 IP Multicast Routing

This chapter describes the components of IP multicast routing, and how to configure IP multicast routing on the BlackDiamond.

For more information on IP multicasting, refer to RFC 1112, RFC 1075, RFC 2236, and other more recent Internet draft documents.

OVERVIEW

IP multicast routing is a function that allows a single IP host to send a packet to a group of IP hosts. This group of hosts can include devices that reside on the local network, within a private network, or outside of the local network.

IP multicast routing consists of the following functions:

- A router that can forward IP multicast packets.
- A router-to-router multicast protocol (for example, Distance Vector Multicast Routing Protocol (DVMRP).
- A method for the IP host to communication its multicast group membership to a router (for example, Internet Group Management Protocol (IGMP)).

DVMRP OVERVIEW

DVMRP is a distance vector protocol that is used to exchange routing and multicast information between routers. Like RIP, DVMRP periodically sends the entire routing table to its neighbors.

DVMRP has a mechanism that allows it to prune and graft multicast trees to reduce the bandwidth consumed by IP multicast traffic.

IGMP OVERVIEW

IGMP is a protocol used by an IP host to register its IP multicast group membership with the router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, a single IP host responds to the query, and group registration is maintained.

IGMP SNOOPING

IGMP snooping adds intelligence to a layer 2 device (such as a switch), to reduce the flooding of IP multicast traffic. The goal of IGMP snooping is to optimize the usage of network bandwidth and prevent multicast traffic from being flooded to parts of the network that do not need to see it. In order to reduce the bandwidth used by IP multicast applications, IGMP snooping is always enabled on the BlackDiamond.

CONFIGURING IP MULTICASTING ROUTING

To configure IP multicast routing, you must do the following:

1 Configure the system for IP unicast routing.

For more information on configuring IP unicast routing, refer to Chapter 8 and Chapter 9.

2 Enable multicast routing on the interface, using the following command:

```
enable ipmcforwarding {vlan <name> | all}
```

3 Enable DVMRP on all IP multicast routing interfaces, using the following command:

```
config dvmrp add {vlan <name> | all}
```

4 Enable DVMRP on the router, using the following command:

```
enable dvmrp
```

Table 10-1 describes the commands used to configure IP multicast routing.

Table 10-1: IP Multicast Routing Configuration Commands

Command	Description
enable dvmrp	Enables DVMRP on the system. The default setting is disabled.
enable ipmcforwarding {<vlan <name> | all>}	Enables IP multicast forwarding on an IP interface. If `all` is specified, all configured IP interfaces are affected. When new IP interfaces are added, `ipmcforwarding` is disabled by default.
config dvmrp add {vlan <name> | all}	Enables DVMRP on an IP interface. When an IP interface is created, DVMRP is enabled by default.
config dvmrp delete {vlan <name> | all}	Disables DVMRP on an IP interface.
config dvmrp vlan <name> timer <probe_interval> <neighbor_timeout_interval>	Configures DVMRP interface timers. Specify the following: ■ `probe_interval` — The amount of time that the system waits between transmitting DVMRP probe messages. The range is 1 to 4,294,967,296 seconds (136 years). The default setting is 10 seconds. ■ `neighbor_timeout_interval` — The amount of time before a DVMRP neighbor route is declared to be down. The range is 1 to 4,294,967,296 seconds (136 years). The default setting is 35 seconds.

Table 10-1: IP Multicast Routing Configuration Commands (continued)

Command	Description
config dvmrp timer <route_report_interval> <route_replacement_time>	Configures the global DVMRP timers. Specify the following: ■ `route_report_interval` — The amount of time the system waits between transmitting periodic route report packets. The range is 1 to 4,294,967,296 seconds (136 years). The default setting is 60 seconds. ■ `route_replacement_time` — The hold-down time before a new route is learned, once the previous route has been deleted. The range is 1 to 4,294,967,296 seconds (136 years). The default setting is 140 seconds.
config ipmc cache timeout <seconds>	Configures the aging time for IP multicast cache entries. The default setting is 300 seconds.

Table 10-2 describes the command used to configure the Internet Gateway Message Protocol (IGMP).

Table 10-2: IGMP Configuration Commands

Command	Description
enable igmp {vlan <name> | all}	Enables IGMP on a router interface. The default setting is enabled.
config igmp <query_interval> <query_response_interval> <last_member_query_interval>	Configures the IGMP timers. Timers are based on RFC2236. Specify the following: ■ `query_interval` — The amount of time, in seconds, the system waits between sending out General Queries. The range is 1 to 4,294,967,296 seconds (136 years). The default setting is 125 seconds. ■ `query_response_interval` — The maximum response time inserted into the periodic General Queries. The range is 1 to 25 seconds. The default setting is 10 seconds. ■ `last_member_query_interval` — The maximum response time inserted into a Group-Specific Query sent in response to a Leave group message. The range is 1 to 25 seconds. The default setting is 1 second.

Table 10-2: IGMP Configuration Commands (continued)

Command	Description
config igmp snooping <router_timeout> <host_timeout>	Configures the IGMP snooping timers. Timers should be set to approximately 2.5 times the router query interval in use on the network. Specify the following: ■ `router_timeout` — The interval, in seconds, between the last time the router was discovered and the current time. The range is 10 to 4,294,967,296 seconds (136 years). The default setting is 260 seconds. ■ `host_timeout` — The interval, in seconds, between the last IGMP group report message from the host and the current time. The range is 10 to 4,294,967,296 seconds (136 years). The default setting is 260 seconds.

CONFIGURATION EXAMPLE

Figure 10-1 is used in Chapter 9 to describe the OSPF configuration on a BlackDiamond. Refer to Chapter 9 for more information about configuring OSPF. In this example, the system labeled IR1 is configured for IP multicast routing.

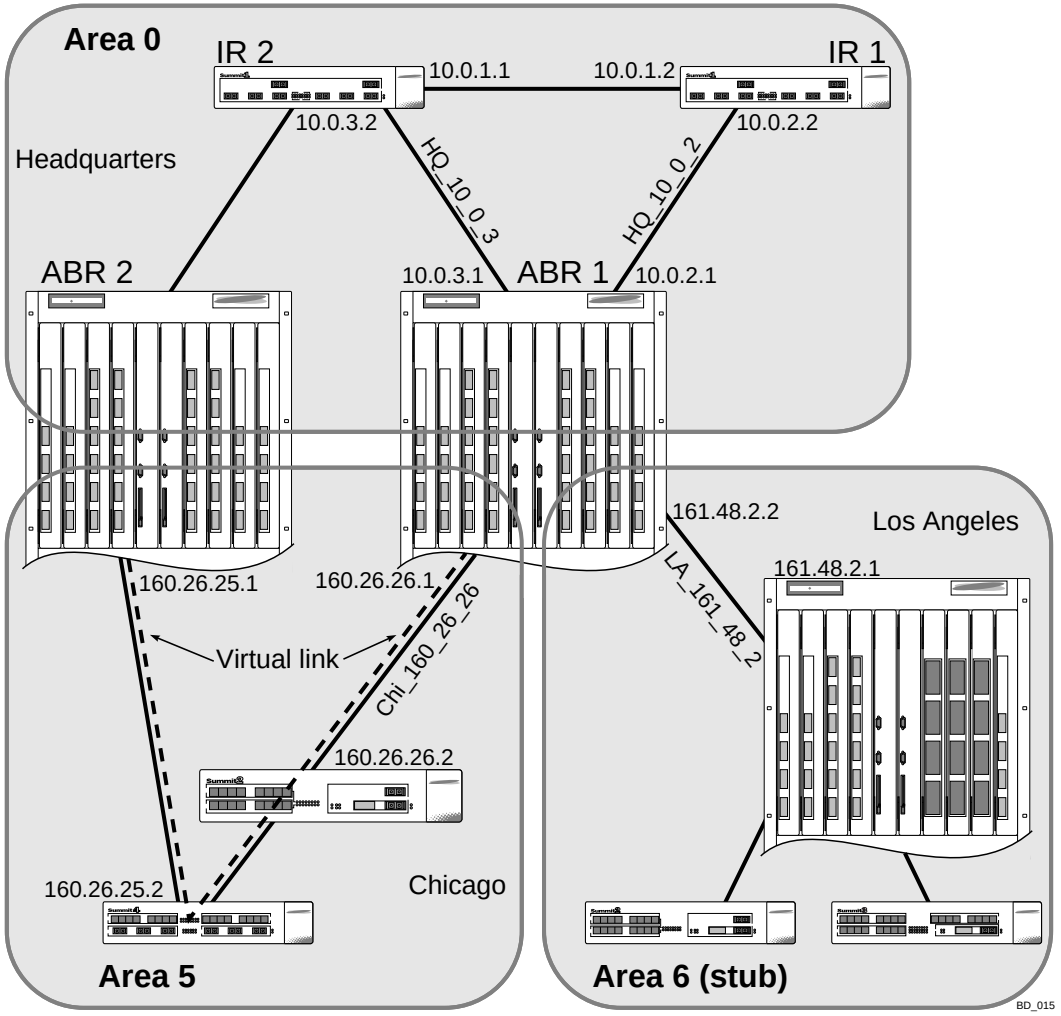

Figure 10-1: IP multicast routing configuration example

CONFIGURATION FOR IR1

The following is the configuration for the router labeled IR1:

```
config vlan HQ_10_0_1 ipaddress 10.0.1.2 255.255.255.0
config vlan HQ_10_0_2 ipaddress 10.0.2.2 255.255.255.0
config ospf add vlan all
enable ipforwarding
enable ospf
enable ipmcforwarding
config dvmrp add vlan all
enable dvmrp
```

DISPLAYING IP MULTICAST ROUTING SETTINGS

To display settings for IP multicast routing components, use the commands listed in Table 10-3.

Table 10-3: IP Multicast Routing Show Commands

Command	Description
show dvmrp {vlan <name> | route | all}	Displays the DVMRP configuration and statistics, or the unicast route table. The default setting is `all`.
show igmp snooping {<vlan <name> | all}	Displays IGMP snooping registration information, and a summary of all IGMP timers and states.
show ipmc cache {<group> {<src_ipaddress> <mask>}} | all}	Displays the IP multicast forwarding cache. Information displayed includes the following: ■ IP group address ■ IP source address and mask ■ Upstream neighbor ■ Interface to upstream neighbor ■ Route expiration timer ■ Routing protocol ■ List of next-hop interfaces and protocols

DELETING AND RESETTING IP MULTICAST SETTINGS

To return IP multicast routing settings to their defaults and disable IP multicast routing functions, use the commands listed in Table 10-4.

Table 10-4: IP Multicast Routing Reset and Disable Commands

Command	Description
disable dvmrp	Disables DVMRP on the system.
disable ipmcforwarding {vlan <name> | all}	Disables IP multicast forwarding.
disable igmp {vlan <name> | all}	Disables IGMP on a router interface.
disable igmp snooping {vlan <name> | all}	Disables IGMP snooping.
unconfig dvmrp [vlan <name> | all]	Resets the DVMRP timers to their default settings.
unconfig igmp	Resets all IGMP settings to their default values and clears the IGMP group table.
clear igmp snooping [vlan <name> | all]	Removes one or more IGMP snooping entries.
clear ipmc cache {<group> {<src_ipaddress> <mask>}} | all]	Resets the IP multicast cache table. If no option is specified, all IP multicast cache entries are flushed.

11 Status Monitoring and Statistics

This chapter describes how to view the current operating status of the BlackDiamond, how to display information in the log, and how to take advantage of available Remote Monitoring (RMON) capabilities.

Viewing statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

STATUS MONITORING

The status monitoring facility provides information about the BlackDiamond. This information may be useful for your technical support representative if you have a problem.

Table 11-1 describes the show commands.

Table 11-1: Switch Monitoring Commands

Command	Description
show account	Displays the account names, access level, number of successful and failed login attempts, and the number of active sessions in the user database. This command is available only to admin level users.
show banner	Displays the user-configured banner.

Table 11-1: Switch Monitoring Commands (continued)

Command	Description
show config	Displays the current configuration to the terminal. You can then capture the output and store it as a file.
show diag	Displays software diagnostics.
show dvmrp {vlan <name> | route | all}	Displays the DVMRP configuration and statistics, or the unicast route table. The default setting is `all`.
show fdb {all | <macaddress> | vlan <name> | <port:portlist> | permanent | qos}	Displays the forwarding database contents including MAC address, associated VLAN, port, age-of-entry configuration method, and status. Providing one of the options acts as a filter on the display. Providing a VLAN name displays all entries for the VLAN. Use the MAC address to locate a specific entry in the FDB.
show gvrp	Displays the current configuration and status of GVRP.
show igmp snooping {<vlan <name> | all}	Displays IGMP snooping registration information, and a summary of all IGMP timers and states.
show iparp {<ip_address> | vlan <name> | all | permanent}	Displays the IP Address Resolution Protocol (ARP) table. You can filter the display by IP address, VLAN, or permanent entries.
show iparp proxy {<ipaddress> {<mask} | all}	Displays the proxy ARP table.
show ipconfig {vlan <name> | all}	Displays configuration information for one or more VLANs, including the following: ■ IP address, subnet mask ■ IP forwarding information ■ BOOTP configuration ■ VLAN name, VLANid ■ Global ICMP configuration ■ Global IGMP configuration ■ Global IRDP configuration
show ipfdb {<ipaddress> | vlan <name> | all}	Displays the contents of the IP forwarding database table.

Table 11-1: Switch Monitoring Commands (continued)

Command	Description
show ipmc cache {<group> {<src_ipaddress> <mask>}} | all}	Displays the IP multicast route table. Information displayed includes the following: ■ IP group address ■ IP source address and mask ■ Upstream neighbor ■ Interface to upstream neighbor ■ Route expiration timer ■ Routing protocol ■ List of next hop interfaces and protocols
show ipqos {<ip_destination_address> <mask> | all}	Displays the IP QoS table.
show iproute vlan {<name> | all | permanent | <ipaddress> <mask>}	Displays the contents of the IP routing table.
show ipstats {vlan [<name> | all]}	Displays statistics of packets handled by the CPU, including the following: ■ inpackets, outpackets ■ ICMP/IGMP statistics ■ IRDP statistics
show log {<priority>} {<subsystem>}	Displays the current snapshot of the log. Options include: ■ `priority` — Filters the log to display message with the selected priority or higher (more critical). Priorities include critical, warning, and informational. If not specified, informational priority messages and higher are displayed. ■ `subsystem` — Filters the log to display messages associated with the selected switch subsystem. Subsystems include Syst, STP Brdg, SNMP, Telnet, VLAN, and Port. If not specified, all subsystems are displayed.
show log config	Displays the log configuration, including the syslog host IP address, the priority level of messages being logged locally, and the priority level of messages being sent to the syslog host.

Table 11-1: Switch Monitoring Commands (continued)

Command	Description
show management	Displays network management configuration and statistics including enable/disable states for Telnet and SNMP, SNMP community strings, authorized SNMP station list, SNMP trap receiver list, and login statistics.
show memory	Displays the current system memory information.
show mirroring	Displays the port-mirroring configuration.
show ospf	Displays global OSPF information.
show ospf area {<areaid> | all}	Displays information about a particular OSPF area, or all OSPF areas.
show ospf interfaces {vlan <name> | area <areaid> | all}	Displays information about one or all OSPF interfaces. If no argument is specific, all OSPF interfaces are displayed.
show ospf lsdb {detail} {area <areaid> | all} {router | network | summary_net | summary_asb | as_external | all}	Displays a table of the current link-state database. You can filter the display using either the area ID or the remote router's router ID, or the link-state ID. The default is `all` with no detail. If detail is specified, each entry includes complete LSA information.
show ospf virtual-link {<areaid> <routerid> | all}	Displays virtual link information about a particular router or all routers.
show port {<slot:portlist>} collisions	Displays real-time collision statistics.
show port {<slot:portlist>} config	Displays the port configuration, including the following: ■ Port state ■ Link state ■ Link speed ■ Duplex mode ■ Flow control ■ Load-sharing information ■ Link media information ■ QoS information

Table 11-1: Switch Monitoring Commands (continued)

Command	Description
show port {<slot:portlist>} information	Displays detailed system-related information, including the following: ■ Port state ■ Link state ■ Autonegotiation state ■ Link speed ■ Duplex mode ■ Load-sharing information ■ VLAN information ■ QoS information
show port {<slot:portlist>} packet	Displays a histogram of packet statistics.
show port {<slot:portlist>} qosmonitor	Displays real-time QoS statistics.
show port {<slot:portlist>} rxerrors	Displays real-time receive error statistics.
show port {<slot:portlist>} stats	Displays real-time port statistics.
show port {<slot:portlist>} txerrors	Displays real-time transmit error statistics.
show port {<slot:portlist>} utilization	Displays real-time port utilization information. Use the [Spacebar] to toggle between packet, byte, and bandwidth utilization information.
show protocol {<protocol> | all}	Displays protocol information including protocol name, protocol fields, and the list of VLANs that use this protocol.
show qosprofile {<qosname> | all}	Displays QoS profile information including the QoS profile name, minimum bandwidth, maximum bandwidth, and priority levels. Also displays the groupings to which this QoS profile is applied.
show rip {vlan <name> | all}	Displays RIP configuration and statistics for one or more VLANs.

Table 11-1: Switch Monitoring Commands (continued)

Command	Description
show rip stat {vlan <name> | all}	Displays RIP-specific statistics. Statistics include the following per interface: ■ Packets transmitted ■ Packets received ■ Bad packets received ■ Bad routes received ■ Number of RIP peers ■ Peer information
show session	Displays the currently active Telnet and console sessions communicating with the switch. Provides the user name, IP address of the incoming Telnet session, whether a console session is currently active, and login time. Sessions are numbered.
show slot <slot>	Displays slot-specific information, including the following: ■ Card type, serial number, part number ■ Current state ■ Port information
show stpd {<stpd_name> | all}	Displays STP information for the one or all STP domains.
show stpd <stpd_name> port <slot:portlist>	Displays port-specific STP information including STP port configuration and state.

Table 11-1: Switch Monitoring Commands (continued)

Command	Description
show switch	Displays the current switch information, including: ■ sysName, sysLocation, sysContact ■ MAC address ■ Current time and time, and system uptime ■ Operating environment (temperature, fans, and power supply status) ■ NVRAM image information (primary/secondary image, date, time, size, version) ■ NVRAM configuration information (primary/secondary configuration, date, time, size, version) ■ Number of MSMs installed ■ MSM configuration (master/slave) ■ Scheduled reboot information ■ 802.1p information ■ System serial number and reworks indicator ■ Software platform ■ System ID ■ Power supply and fan status
show version	Displays the hardware and software versions currently running on the switch. Displays the switch serial number and version numbers of MSM and I/O modules.
show vlan {<name> | all}	When used with the keyword `all`, or with no named VLANs, displays a summary list of VLAN names with a portlist and associated status of each. When used with a named identifier, displays port information including membership list, IP address, tag information.

SLOT DIAGNOSTICS

The BlackDiamond system provides a facility for running normal or extended diagnostics on an I/O module or a Management Switch Fabric Module (MSM) without affecting the operation of the rest of the system.

If you select to run the diagnostic routine on an I/O module, that module is taken off-line while the diagnostic test is performed. Traffic to and from the ports on the module are temporarily unavailable. Once the diagnostic test is completed, the I/O module is reset and becomes operational again.

You can run normal or extended diagnostics on the slave MSM. The normal diagnostic routing is a short series of tests that do not test all the internal ASIC functions. The extended diagnostic routine tests coverage of all MSM components including the internal ASIC functions. The slave MSM is taken off-line while the diagnostic test is performed. It is reset and operational once the test is completed.

If you want to run the diagnostic routine on the master MSM, you must set it in diagnostic mode and then reboot the switch. When you reboot, the master MSM becomes the slave MSM (and vice versa) so that the diagnostic routing can be performed.

To run the diagnostic routine, use the command

```
run diag [normal | extended] [<slot> | a | b]
```

where the following is true:

- `[normal | extended]` — selects the diagnostic routine
- `<slot>` — specifies the slot number of an I/O module
- `a | b` — specifies the slot letter of an MSM

PORT STATISTICS

The BlackDiamond provides a facility for viewing port statistic information. The summary information lists values for the current counter against each port on each operational module in the system, and it is refreshed approximately every 2 seconds. Values are displayed to nine digits of accuracy.

To view port statistics, use the following command:

```
show port <slot:portlist> stats
```

The following port statistic information is collected by the switch:

- **Link Status —** The current status of the link. Options are
  - **—** Ready — The port is ready to accept a link.
  - **—** Active — The link is present at this port.
  - **—** Chassis — The link is connected to a Summit Virtual Chassis.
- **Transmit Packet Count (Tx Pkt Count) —** The number of packets that have been successfully transmitted by the port.
- **Transmit Byte Count (Tx Byte Count) —** The total number of data bytes successfully transmitted by the port.
- **Total Collisions —** The total number of collisions seen by the port, regardless of whether a device connected to the port participated in any of the collisions.
- **Received Packet Count (Rx Pkt Count) —** The total number of good packets that have been received by the port.
- **Received Byte Count (RX Byte Count) —** The total number of bytes that were received by the port, including bad or lost frames. This number includes bytes contained in the Frame Check Sequence (FCS), but excludes bytes in the preamble.
- **Receive Broadcast (RX Bcast)** — The total number of frames received by the port that are addressed to a broadcast address.
- **Receive Multicast (RX Mcast)** — The total number of frames received by the port that are addressed to a multicast address.

PORT ERRORS

The BlackDiamond keeps track of errors for each port.

To view port transmit errors, use the following command:

```
show port <slot:portlist> txerrors
```

The following port transmit error information is collected by the system:

- **Link Status —** The current status of the link. Options are
  - **—** Ready — The port is ready to accept a link.
  - **—** Active — The link is present at this port.
- **Transmit Collisions (TX Coll) —** The total number of collisions seen by the port, regardless of whether a device connected to the port participated in any of the collisions.
- **Transmit Late Collisions (TX Late) —** The total number of collisions that have occurred after the port's transmit window has expired.
- **Transmit Deferred Frames (TX Def) —** The total number of frames that were transmitted by the port after the first transmission attempt was deferred by other network traffic.
- **Transmit Errored Frames (TX Err) —** The total number of frames that were not completely transmitted by the port because of network errors (such as late collisions or excessive collisions).

To view port receive errors, use the following command:

```
show port <slot:portlist> rxerrors
```

The following port receive error information is collected by the switch:

- **Receive Bad CRC Frames (RX CRC) —** The total number of frames received by the port that were of the correct length, but contained a bad FCS value.
- **Receive Oversize Frames (RX Over) —** The total number of good frames received by the port that were of greater than the supported maximum length of 1,522 bytes.
- **Receive Undersize Frames (RX Under) —** The total number of frames received by the port that were less than 64 bytes long.
- **Receive Jabber Frames (RX Jab) —** The total number of frames received by the port that was of greater than the support maximum length and had a Cyclic Redundancy Check (CRC) error.

- **Receive Alignment Errors (RX Align) —** The total number of frames received by the port that occurs if a frame has a CRC error and does not contain an integral number of octets.
- **Receive Frames Lost (RX Lost) —** The total number of frames received by the port that were lost because of buffer overflow in the switch.

PORT MONITORING DISPLAY KEYS

Table 11-2 describes the keys used to control the displays that appear when you issue any of the `show port` commands.

Table 11-2: Port Monitoring Display Keys

Key(s)	Description
U	Displays the previous page of ports.
D	Displays the next page of ports.
[Esc] or [Return]	Exits from the screen.
0	Clears all counters.
[Space]	Cycles through the following screens: ■ Packets per second ■ Bytes per second ■ Percentage of bandwidth Available using the `show port utilization` command only.

LOGGING

The BlackDiamond log tracks all configuration and fault information pertaining to the device. Each entry in the log contains the following information:

- **Timestamp** — The timestamp records the month and day of the event, along with the time (hours, minutes, and seconds) in the form HH:MM:SS. If the event was caused by a user, the user name is also provided.

- **Fault level** — Table 11-3 describes the three levels of importance that the system can assign to a fault.

Table 11-3: Fault Levels Assigned by the System

Level	Description
Critical	A desired switch function is inoperable. The switch may need to be reset.
Warning	A noncritical error that may lead to a function failure.
Informational	Actions and events that are consistent with expected behavior.

- **Subsystem** — The facility refers to the specific functional area to which the error refers. Table 11-4 describes the subsystems.

Table 11-4: Fault Log Subsystems

Subsystem	Description
Syst	General system-related information. Examples include memory, power supply, security violations, fan failure, overheat condition, and configuration mode.
STP	STP information. Examples include an STP state change.
Brdg	Bridge-related functionality. Examples include low table space and queue overflow.
SNMP	SNMP information. Examples include community string violations.
Telnet	Information related to Telnet login and configuration performed by way of a Telnet session.
VLAN	VLAN-related configuration information.
Port	Port management-related configuration. Examples include port statistics and errors.

- **Message** — The message contains the log information with text that is specific to the problem.

LOCAL LOGGING

The BlackDiamond maintains 1,000 messages in its internal log. You can display a snapshot of the log at any time by using the command

```
show log {<priority>} {<subsystem>}
```

where the following is true:

- `priority` — Filters the log to display message with the selected priority or higher (more critical). Priorities include critical, warning, and informational. If not specified, informational priority messages and higher are displayed.
- `subsystem` — Filters the log to display messages associated with the selected switch subsystem. Subsystems include Syst, STP Brdg, SNMP, Telnet, VLAN, and Port. If not specified, all subsystems are displayed.

REAL-TIME DISPLAY

In addition to viewing a snapshot of the log, you can configure the system to maintain a running real-time display of log messages on the console. To turn on the log display, enter the following command:

```
enable log display
```

To configure the log display, use the following command:

```
config log display {<priority>} {<subsystem>}
```

If `priority` is not specified, only messages of critical priority are displayed. If the `subsystem` is not specified, all subsystems are displayed.

If you enable the log display on a terminal connected to the console port, your settings will remain in effect even after your console session is ended (unless you explicitly disable the log display).

When using a Telnet connection, if your Telnet session is disconnected (because of the inactivity timer, or for other reasons), the log display is automatically halted. You must restart the log display by using the `enable log display` command.

REMOTE LOGGING

In addition to maintaining an internal log, the BlackDiamond supports remote logging by way of the UNIX syslog host facility. To enable remote logging, do the following:

- Configure the syslog host to accept and log messages.
- Enable remote logging by using the following command:

  ```
  enable syslog
  ```

- Configure remote logging by using the following command:

  ```
  config syslog <ipaddress> <facility> {<priority>} {<subsystem>}
  ```

Specify the following:

- `ipaddress` — The IP address of the syslog host.
- `facility` — The syslog facility level for local use. Options include `local0` through `local7`.
- `priority` — Filters the log to display message with the selected priority or higher (more critical). Priorities include critical, warning, and informational. If not specified, only critical priority messages are sent to the syslog host.
- `subsystem` — Filters the log to display messages associated with the selected subsystem. Subsystems include Syst, STP Brdg, SNMP, Telnet, VLAN, and Port. If not specified, all subsystems are sent to the syslog host.

Refer to your UNIX documentation for more information about the syslog host facility.

LOGGING COMMANDS

The commands described in Table 11-5 allow you to configure logging options, reset logging options, display the log, and clear the log.

Table 11-5: Logging Commands

Command	Description
config log display {<priority>} {<subsystem>}	Configures the real-time log display. Options include: ■ `priority` — Filters the log to display messages with the selected priority or higher (more critical). Priorities include critical, warning, and informational. If not specified, informational priority messages and higher are displayed. ■ `subsystem` — Filters the log to display messages associated with the selected switch subsystem. Subsystems include Syst, STP Brdg, SNMP, Telnet, VLAN, and Port. If not specified, all subsystems are displayed.
config syslog <ip_address> <facility> {<priority>} {<subsystem>}	Configures the syslog host address and filter messages sent to the syslog host. Options include: ■ `ipaddress` — The IP address of the syslog host. ■ `facility` — The syslog facility level for local use. ■ `priority` — Filters the log to display messages with the selected priority or higher (more critical). Priorities include critical, warning, and informational. If not specified, only critical priority messages and are sent to the syslog host. ■ `subsystem` — Filters the log to display messages associated with the selected switch subsystem. Subsystems include Syst, STP Brdg, SNMP, Telnet, VLAN, and Port. If not specified, all subsystems are sent to the syslog host.
enable log display	Enables the log display.
enable syslog	Enables logging to a remote syslog host.
disable log display	Disables the log display.
disable syslog	Disables logging to a remote syslog host.

Table 11-5: Logging Commands (continued)

Command	Description
show log {<priority>} {<subsystem>}	Displays the current snapshot of the log. Options include: ■ `priority` — Filters the log to display message with the selected priority or higher (more critical). Priorities include critical, warning, and informational. If not specified, informational priority messages and higher are displayed. ■ `subsystem` — Filters the log to display messages associated with the selected switch subsystem. Subsystems include Syst, STP Brdg, SNMP, Telnet, VLAN, and Port. If not specified, all subsystems are displayed.
show log config	Displays the log configuration, including the syslog host IP address, the priority level of messages being logged locally, and the priority level of messages being sent to the syslog host.
clear counters	Clears all switch statistics and port counters.
clear log {static}	Clears the log. If `static` is specified, the critical log messages are also cleared.

RMON

Using the Remote Monitoring (RMON) capabilities of the BlackDiamond allows network administrators to improve system efficiency and reduce the load on the network.

The following sections explain more about the RMON concept and the RMON features supported by the BlackDiamond.

You can only use the RMON features of the system if you have an RMON management application.

ABOUT RMON

RMON is the common abbreviation for the Remote Monitoring Management Information Base (MIB) system defined by the Internet Engineering Task Force (IETF) documents RFC 1271 and RFC 1757, which allows you to monitor LANs remotely.

A typical RMON setup consists of the following two components:

- **RMON probe** — An intelligent, remotely controlled device or software agent that continually collects statistics about a LAN segment or VLAN. The probe transfers the information to a management workstation on request, or when a predefined threshold is crossed.
- **Management workstation** — Communicates with the RMON probe and collects the statistics from it. The workstation does not have to be on the same network as the probe, and can manage the probe by in-band or out-of-band connections.

RMON FEATURES OF THE BLACKDIAMOND

The IETF defines nine groups of Ethernet RMON statistics. The BlackDiamond supports the following four of these groups:

- Statistics
- History
- Alarms
- Events

This section describes these groups, and discusses how they can be used.

STATISTICS

The RMON Ethernet Statistics group provides traffic and error statistics showing packets, bytes, broadcasts, multicasts, and errors on a LAN segment or VLAN.

Information from the Statistics group is used to detect changes in traffic and error patterns in critical areas of the network.

HISTORY

The History group provides historical views of network performance by taking periodic samples of the counters supplied by the Statistics group. The group features user-defined sample intervals and bucket counters for complete customization of trend analysis.

The group is useful for analysis of traffic patterns and trends on a LAN segment or VLAN, and to establish baseline information indicating normal operating parameters.

ALARMS

The Alarms group provides a versatile, general mechanism for setting threshold and sampling intervals to generate events on any RMON variable. Both rising and falling thresholds are supported, and thresholds can be on the absolute value of a variable or its delta value. In addition, alarm thresholds may be autocalibrated or set manually.

Alarms inform you of a network performance problem and can trigger automated action responses through the Events group.

EVENTS

The Events group creates entries in an event log and/or sends SNMP traps to the management workstation. An event is triggered by an RMON alarm. The action taken can be configured to ignore it, to log the event, to send an SNMP trap to the receivers listed in the trap receiver table, or to both log and send a trap. The RMON traps are defined in RFC 1757 for rising and falling thresholds.

Effective use of the Events group saves you time. Rather than having to watch real-time graphs for important occurrences, you can depend on the Event group for notification. Through the SNMP traps, events can trigger other actions, providing a mechanism for an automated response to certain occurrences.

RMON AND THE BLACKDIAMOND

RMON requires one probe per LAN segment, and standalone RMON probes have traditionally been expensive. Therefore, Extreme's approach has been to build an inexpensive RMON probe into the agent of each system. This allows RMON to be widely deployed around the network without costing more than traditional network management. The BlackDiamond accurately maintains RMON statistics at the maximum line rate of all of its ports.

For example, statistics can be related to individual ports. Also, because a probe must be able to see all traffic, a stand-alone probe must be attached to a nonsecure port. Implementing RMON in the switch means that all ports can have security features enabled.

EVENT ACTIONS

The actions that you can define for each alarm are shown in Table 11-6.

Table 11-6: Event Actions

Action	High Threshold
No action	
Notify only	Send trap to all trap receivers.
Notify and log	Send trap; place entry in RMON log

To be notified of events using SNMP traps, you must configure one or more trap receivers, as described in Chapter 2.

12 Using ExtremeWare Vista

ExtremeWare Vista is device-management software running in the BlackDiamond that allows you to access the system over a TCP/IP network, using a standard Web browser. Any properly configured standard Web browser that supports frames and JavaScript (such as Netscape Navigator 3.0 or higher, or Microsoft Internet Explorer 3.0 or higher) can be used to manage the system.

ExtremeWare Vista provides a subset of the command-line interface commands available for configuring and monitoring the BlackDiamond. If a particular command is not available using ExtremeWare Vista, you must use the command-line interface to access the desired functionality.

Enabling and Disabling Web Access

By default, Web access is enabled on the BlackDiamond. To disable it, use the following command:

```
disable web
```

To re-enable Web access, use the following command:

```
enable web
```

You will need to reboot the system in order for these changes to take effect.

For more information on rebooting, refer to Chapter 13.

To use ExtremeWare Vista, at least one VLAN must be assigned an IP address.

For more information on assigning an IP address, refer to Chapter 2.

SETTING UP YOUR BROWSER

In general, the default settings that come configured on your browser work well with ExtremeWare Vista. The following are recommended settings that you can use to improve the display features and functionality of ExtremeWare Vista:

- After downloading a newer version of the BlackDiamond image, clear the browser disk and memory cache to see the updated menu screens. You must clear the cache while at the main ExtremeWare Vista Logon screen, so that all underlying .GIF files are updated.
- Check for newer versions of stored pages. Every visit to the page should be selected as a cache setting.

  If you are using Netscape Navigator, configure the cache option to check for changes "Every Time" you request a page.

  If you are using Microsoft Internet Explorer, configure the Temporary Internet Files setting to check for newer versions of stored pages by selecting "Every visit to the page."
- Images must be auto-loaded.
- Use a high-resolution monitor to maximize the amount of information displayed in the content frame. The recommended resolution is 1024 x 768 pixels. You can also use 800 x 600 pixels.
- Turn off one or more of the browser toolbars to maximize the viewing space of the ExtremeWare Vista content screen.
- If you will be using ExtremeWare Vista to send an e-mail to the Extreme Networks Technical Support department, configure the e-mail settings in your browser.
- Configure the browser to use the following recommended fonts:

  — Proportional font—Times New Roman

  — Fixed-width font—Courier New

ACCESSING EXTREMEWARE VISTA

To access the default home page of the BlackDiamond, enter the following URL in your browser:

```
http://<ip_address>
```

When you access the home page of the system, you are presented with the Login screen. Enter your user name and password in the appropriate fields, and click OK.

If you have entered the name and password of an administrator-level account, you have access to all ExtremeWare Vista pages. If you have used a user-level account name and password, you only have access to the Statistics and Support information.

For more information on assigning user names, levels, and passwords, refer to Chapter 2.

If multiple people access the same switch using ExtremeWare Vista, you might see the following error message:

```
Web:server busy
```

To correct this situation, log out of the switch and log in again.

NAVIGATING EXTREMEWARE VISTA

After logging in to the switch, the ExtremeWare Vista home page is displayed.

ExtremeWare Vista divides the browser screen into the following sections:

- Task frame
- Content frame
- Standalone buttons

TASK FRAME

The task frame has two sections: menu buttons submenu links. There are four task buttons, as follows:

- Configuration
- Statistics
- Support
- Logout

Below the task buttons are options. Options are specific to the task button that you select. When you select an option, the information displayed in the content frame changes. However, when you select a new task button, the content frame does not change until you select a new option.

CONTENT FRAME

The content frame contains the main body of information in ExtremeWare Vista. For example, if you select an option from the Configuration task button, enter configuration parameters in the content frame. If you select the Statistics task button, statistics are displayed in the content frame.

BROWSER CONTROLS

Browser controls include drop-down list boxes, check boxes, and multi-select list boxes. A multi-select list box has a scrollbar on the right side of the box. Using a multi-select list box, you can select a single item, all items, a set of contiguous items, or multiple non-contiguous items. Table 12-1 describes how to make selections from a multi-select list box.

Table 12-1: Multi-Select List Box Key Definitions

Selection Type	Key Sequence
Single item	Click the item using the mouse.
All items	Click the first item, and drag to the last item.

Table 12-1: Multi-Select List Box Key Definitions (continued)

Selection Type	Key Sequence
Contiguous items	Click the first desired item, and drag to the last desired item.
Selected non-contiguous items	Hold down [Control], click the first desired item, click the next desired item, and so on.

STATUS MESSAGES

Status messages are displayed at the top of the content frame. There are four types of status messages, as follows:

- Information—Displays information that is useful to know prior to, or as a result of, changing configuration options.
- Warning—Displays warnings about the switch configuration.
- Error—Displays errors caused by incorrectly configured settings.
- Success—Displays informational messages after you click Submit. The message displayed reads, “Request was submitted successfully.”

STANDALONE BUTTONS

At the bottom of some of the content frames is a section that contains standalone buttons. Standalone buttons are used to perform tasks that are not associated with a particular configuration option. An example of this is the Reboot Switch button.

SAVING CHANGES

There are two ways to save your changes to non-volatile RAM (NVRAM) storage using ExtremeWare Vista:

- Select Save Configuration from the Configuration task button, Switch option.

  This field contains a drop-down list box that allows you to select either the primary or secondary configuration area. After you select the configuration area, click Submit to save the changes.

For more information on the primary and secondary configuration areas, refer to Chapter 13.

- Click the Logout button.

  If you attempt to log out without saving your changes, ExtremeWare Vista prompts you to save your changes.

  If you select Yes, the changes are saved to the selected configuration area. To change the selected configuration area, you must go to the Configuration task button, Switch option.

FILTERING INFORMATION

Some pages have a Filter button. The Filter button is used to display a subset of information on a given page. For example, on the OSPF configuration page, you can configure authentication based on the VLAN, area identifier, or virtual link. Once you select a filtering option and click the Filter button, the form that provides the configuration options displays the available interfaces in the drop down menu, based on your filtering selection.

Similarly, in certain Configuration and Statistics pages, information is shown based on a particular slot.

Because the BlackDiamond allows you to preconfigure modules without having them physically available in the chassis, the configuration pages offer a drop down menu to select any module card that has been configured on the system, whether or not the module is physically available. By default, information for the first configured module that is found in the chassis is displayed on the page. You can configure available slots and ports by filtering on a selected module from the Sort by Slot drop down menu.

On the Statistics pages, you can only view information for cards that are configured and physically inserted into the BlackDiamond chassis. On these pages, the Sort by Slot drop down menu displays only these modules.

DO A GET WHEN CONFIGURING A VLAN

When configuring a VLAN using ExtremeWare Vista, prior to editing the VLAN configuration you must first click on the `get` button to ensure that subsequent edits are applied to the correct VLAN. If you do not click on the `get` button and you submit the changes, the changes will be made to the VLAN that was previously displayed.

If you configure a VLAN and then delete it, the *default* VLAN is shown in the VLAN name window, but the VLAN information contained in the lower portion of the page is not updated. Click on the get button to update the display.

SENDING SCREEN OUTPUT TO EXTREME NETWORKS

If Extreme Networks requests that you e-mail the output of a particular ExtremeWare Vista screen, do the following:

1. Click on the content frame of the screen that you must send.
2. From Netscape Navigator, select Save Frame As from the File menu, and enter a name for the file.
3. From Microsoft Internet Explorer 3.0, select Save As File from the File menu, and enter a name for the file.
4. From Microsoft Internet Explorer 4.0, right-click in the content frame, select View Source, and save the HTML text by copying it and pasting it into a text editor.
5. Attach the file to the e-mail message that you are sending to Extreme Networks.

13 Software Upgrade and Boot Options

This chapter describes the procedure for upgrading the system software image. This chapter also discusses how to save and load a primary and secondary image and configuration file on the BlackDiamond.

Downloading a New Image

The image file contains the executable code that runs on the BlackDiamond. It comes preinstalled from the factory. As new versions of the image are released, you should upgrade the software running on your system.

The image is upgraded by using a download procedure from either a Trivial File Transfer Protocol (TFTP) server on the network or from a PC connected to the serial port using the XMODEM protocol. Downloading a new image involves the following steps:

- Load the new image onto a TFTP server on your network, if you will be using TFTP.
- Load the new image onto a PC, if you will be using XMODEM.

- Download the new image to the BlackDiamond using the command

  `download image [xmodem | <ipaddress> <filename>] {primary | secondary}`

  where the following is true:

  `xmodem` — Indicates that you will be using XMODEM over the serial port.

  `ipaddress` — Is the IP address of the TFTP server.

  `filename` — Is the filename of the new image.

  `primary` — Indicates the primary image.

  `secondary` — Indicates the secondary image.

The BlackDiamond can store up to two images: a primary and a secondary. When you download a new image, you must select into which image space (primary or secondary) you want the new image to be placed. If you have two MSMs installed in the BlackDiamond, the downloaded image is saved to the same location on each one.

You can select which image the BlackDiamond will load on the next reboot by using the following command:

`use image {primary | secondary}`

If you do not specify which image to use, the system automatically loads the primary image.

REBOOTING THE SYSTEM

To reboot the system, use the following command:

`reboot {<date> <time> | cancel}`

where `date` is the date and `time` is the time (using a 24-hour clock format) when the switch will be rebooted. The values use the following format:

`mm/dd/yyyy hh:mm:ss`

If you do not specify a reboot time, the reboot happens immediately following the command, and any previously scheduled reboots are cancelled. To cancel a previously scheduled reboot, use the `cancel` option.

SAVING CONFIGURATION CHANGES

The configuration is the customized set of parameters that you have selected to run on the BlackDiamond. As you make configuration changes, the new settings are stored in run-time memory. Settings that are stored in run-time memory are not retained by the switch when the switch is rebooted. To retain the settings, and have them be loaded when you reboot the switch, you must save the configuration to NVRAM.

The BlackDiamond can store two different configurations: a primary and a secondary. When you save configuration changes, you can select to which configuration you want the changes saved. If you do not specify, the changes are saved to the configuration area currently in use.

If you have made a mistake, or you must revert to the configuration as it was before you started making changes, you can tell the system to use the secondary configuration on the next reboot.

To save the configuration, use the following command:

```
save {config} {primary | secondary}
```

To use the configuration, use the following command:

```
use config {primary | secondary}
```

The configuration takes effect on the next reboot.

If the system is rebooted while in the middle of a configuration save, the switch boots to factory default settings. The configuration that is not in the process of being saved is unaffected.

RETURNING TO FACTORY DEFAULTS

To return the BlackDiamond configuration to factory defaults, use the following command:

```
unconfig switch
```

This command resets the entire configuration, with the exception of user accounts and passwords that have been configured.

To reset all parameters, use the following command:

```
unconfig switch all
```

USING TFTP TO UPLOAD THE CONFIGURATION

You can upload the current configuration to a TFTP server on your network. The uploaded ASCII file retains the command-line interface format. This allows you to do the following:

- Modify the configuration using a text editor, and later download a copy of the file to the same system, or to one or more different systems.
- Send a copy of the configuration file to Extreme Networks Technical Support for problem-solving purposes.
- Automatically upload the configuration file every day, so that the TFTP server can archive the configuration on a daily basis. Because the filename is not changed, the configured file stored in the TFTP server is overwritten every day.

To upload the configuration, use the command

```
upload config <ipaddress> <filename> {every <time> | cancel}
```

where the following is true:

- `ipaddress` — Is the IP address of the TFTP server.
- `filename` — Is the name of the ASCII file. The filename can be up to 255 characters long, and can not include any spaces, commas, quotation marks, or special characters.
- `every <time>` — Specifies the time of day you want the configuration automatically uploaded on a daily basis.
- `cancel` — Cancels automatic upload, if it has been previously configured.

USING TFTP TO DOWNLOAD THE CONFIGURATION

You can download a previously saved configuration from a TFTP server. To download a configuration, use the following command:

```
download config <ipaddress> <filename>
```

After the ASCII configuration file is downloaded by way of TFTP, you are prompted to reboot the system. The downloaded configuration file is stored in an area of system memory, and is not retained if the BlackDiamond has a power failure. When the system is rebooted, it treats the downloaded configuration file as a script of command-line interface commands. After the script is executed, you are prompted to save the configuration.

BOOT OPTION COMMANDS

Table 13-1 lists the commands associated with BlackDiamond boot options.

Table 13-1: Boot Option Commands

Command	Description
show config	Displays the current configuration to the terminal. You can then capture the output and store it as a file.
download bootrom <ipaddress> <filename>	Downloads a BOOT ROM image from a TFTP server. The downloaded image replaces the BOOT ROM in the onboard FLASH memory.
download config <ipaddress> <filename>	Downloads a previously saved ASCII configuration file from a specific IP host.
download image [xmodem | <ipaddress> <filename>] {primary | secondary}	Downloads a new image by way of XMODEM using the serial port, or from a TFTP server over then network. If no parameters are specified, the image is saved to the current image.
copy image from [msm-a | msm-b] to [msm-a | msm-b]	Copies the primary and secondary image stored in the onboard FLASH memory from one MSM to another MSM.

Table 13-1: Boot Option Commands (continued)

Command	Description
copy configuration from [msm-a | msm-b] to [msm-a | msm-b]	Copies the primary and secondary configuration stored in the onboard FLASH memory from one MSM to another MSM. To copy the runtime configuration, use the `save configuration` command.
reboot {<date> <time> | cancel}	Reboots the system at the date and time specified. If you do not specify a reboot time, the reboot happens immediately following the command, and any previously scheduled reboots are cancelled. To cancel a previously scheduled reboot, use the `cancel` option.
save {config} {primary | secondary}	Saves the current configuration to NVRAM. You can specify the primary or secondary configuration area. If not specified, the configuration is saved to the configuration area currently in use.
upload config <ipaddress> <filename> {every <time> | cancel}	Uploads the current runtime configuration to the specified TFTP server. If `every <time>` is specified, the switch automatically saves the configuration to the server once per day, at the specified time. To cancel automatic upload, use the `cancel` option. If no options are specified, the current configuration is uploaded immediately.
use config {primary | secondary}	Configures the system to use a particular configuration on the next reboot. Options include the primary configuration area, or the secondary configuration area. If not specified, the primary configuration area is used.
use image [primary | secondary]	Configures the system to use a particular image on the next reboot.

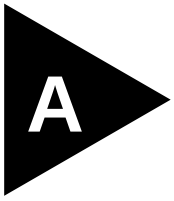

Supported Standards

The following is a list of software standards supported on the BlackDiamond system.

SNMP	**Terminal Emulation**
SNMP protocol (RFC 1157)	Telnet (RFC 854)
MIB-II (RFC 1213)	HTTP 1.0
Bridge MIB (RFC 1493)	**Protocols Used for Administration**
Interfaces MIB (RFC 1573)	UDP (RFC 768)
RMON MIB (RFC 1757)	IP (RFC 791)
802.3 MAU MIB (RFC 2239)	ICMP (RFC 792)
IP Forwarding MIB (RFC 1354)	TCP (RFC 793)
OSPF2 MIB (RFC 1850)	ARP (RFC 826)
RIP2 MIB (RFC 1724)	TFTP (RFC 783)
	BOOTP (RFC 1271)

B Troubleshooting

If you encounter problems when using the BlackDiamond, this appendix may be helpful. If you have a problem not listed here or in the “Release Notes”, contact your local technical support representative.

LEDs

Power LED does not light:

Check that the power cable is firmly connected to the device and to the supply outlet.

On powering-up, the MGMT LED lights yellow:

The device has failed its Power On Self Test (POST) and you should contact your supplier for advice.

A link is connected, but the Status LED does not light:

Check that

- All connections are secure.
- Cables are free from damage.
- The devices at both ends of the link are powered-up.

- Both ends of the Gigabit link are set to the same autonegotiation state.

  Both sides if the Gigabit link must be enabled or disabled. It the two are different, typically the side with autonegotiation disabled will have the link LED list, and the side with autonegotiation enabled will not list. The default configuration for a Gigabit port is autonegotiation enabled. This can be verified by entering the following command:

  ```
  show port config
  ```

On power-on, some I/O modules do not boot:

Check if you are using 110V power input. the BlackDiamond powers only up to four modules if it is connected to a 110V outlet.

Error LED on the MSM turns amber:

Check the syslog message for a “critical” software error.

Status LED on the I/O module turns amber:

Check the syslog message for a related I/O module error. If the error is an inserted an I/O module that conflicts with the software configuration, use one of the following commands to reset the slot configuration:

```
clear slot
```

```
config slot <slot> module [f32t | f32f | g4x | g6x]
```

Otherwise, contact Extreme Networks for further assistance.

USING THE COMMAND-LINE INTERFACE

The initial welcome prompt does not display:

Check that your terminal or terminal emulator is correctly configured.

For console port access, you may need to press [Return] several times before the welcome prompt appears.

Check the settings on your terminal or terminal emulator. The settings are 9600 baud, 8 data bits, 1 stop bit, no parity, XON/OFF flow control enabled.

The SNMP Network Manager cannot access the device:

Check that the device IP address, subnet mask, and default router are correctly configured, and that the device has been reset.

Check that the device IP address is correctly recorded by the SNMP Network Manager (refer to the user documentation for the Network Manager).

Check that the community strings configured for the system and Network Manager are the same.

Check that SNMP access was not disabled for the system.

The Telnet workstation cannot access the device:

Check that the device IP address, subnet mask and default router are correctly configured, and that the device has been reset. Ensure that you enter the IP address of the switch correctly when invoking the Telnet facility. Check that Telnet access was not disabled for the switch. If you attempt to log in and the maximum number of Telnet sessions are being used, you should receive an error message indicating so.

Traps are not received by the SNMP Network Manager:

Check that the SNMP Network Manager's IP address and community string are correctly configured, and that the IP address of the Trap Receiver is configured properly on the system.

The SNMP Network Manager or Telnet workstation can no longer access the device:

Check that Telnet access or SNMP access is enabled.

Check that the port through which you are trying to access the device has not been disabled. If it is enabled, check the connections and network cabling at the port.

Check that the port through which you are trying to access the device is in a correctly configured VLAN.

Try accessing the device through a different port. If you can now access the device, a problem with the original port is indicated. Re-examine the connections and cabling.

A network problem may be preventing you accessing the device over the network. Try accessing the device through the console port.

Check that the community strings configured for the device and the Network Manager are the same.

Check that SNMP access was not disabled for the system.

Permanent entries remain in the FDB:

If you have made a permanent entry in the FDB (which requires you to specify the VLAN to which it belongs and then delete the VLAN), the FDB entry will remain. Though causing no harm, you must manually delete the entry from the FDB if you want to remove it.

Default and Static Routes:

If you have defined static or default routes, those routes will remain in the configuration independent of whether the VLAN and VLAN IP address that used them remains. You should manually delete the routes if no VLAN IP address is capable of using them.

You forget your password and cannot log in:

If you are not an administrator, another user having administrator access level can log in, delete your user name, and create a new user name for you, with a new password.

Alternatively, another user having administrator access level can log in and initialize the device. This will return all configuration information (including passwords) to the initial values.

In the case where no one knows a password for an administrator level user, contact your supplier.

VLANs

You cannot add a port to a VLAN:

If you attempt to add a port to a VLAN and get an error message similar to

```
localhost:7 # config vlan marketing add port 1:1,1:2
ERROR: Protocol conflict on port 1:5
```

you already have a VLAN using untagged traffic on a port. Only one VLAN using untagged traffic can be configured on a single physical port.

VLAN configuration can be verified by using the following command:

```
show vlan <name>
```

The solution for this error is to remove ports 1 and 2 from the VLAN currently using untagged traffic on those ports. If this were the "default" VLAN, the command would be

```
localhost:23 # config vlan default del port 1:1,1:2
```

which should now allow you to re-enter the previous command without error as follows:

```
localhost:26 # config vlan red add port 1:1,1:2
```

VLAN names:

There are restrictions on VLAN names. They cannot contain white spaces and cannot start with a numeric value unless you use quotation marks around the name. If a name contains white spaces, starts with a numeric, or contains non-alphabetical characters, you must use quotation marks whenever referring to the VLAN name.

802.1Q links do not work correctly:

Remember that VLAN names are only locally significant through the command-line interface. For two switches to communicate across a 802.1Q link, the VLAN ID for the VLAN on one switch should have a corresponding VLAN ID for the VLAN on the other switch.

If you are connecting to a third-party device and have checked that the VLAN IDs are the same, the Ethertype field used to identify packets as 802.1Q packets may differ between the devices. The default value used by the switch is **8100**. If the third-party device differs from this and cannot be changed, you may change the 802.1Q Ethertype used by the BlackDiamond with the following command:

```
config dot1p ethertype <ethertype>
```

Changing this parameter changes how the system recognizes all tagged frames received, as well as the value it inserts in all tagged frames it transmits.

VLANs, IP Addresses and default routes:

The system can have an IP address for each configured VLAN. It is only necessary to have an IP address associated with a VLAN if you intend to manage (Telnet, SNMP, ping) through that VLAN. You can also configure multiple default routes for the system. The system first tries the default route with the lowest cost metric.

STP

You have connected an endstation directly to the switch and the endstation fails to boot correctly:

The BlackDiamond has STP enabled, and the endstation is booting before the STP initialization process is complete. Specify that STP has been disabled for that VLAN, or turn off STP for the switch ports of the endstation and devices to which it is attempting to connect, and then reboot the endstation.

The switch keeps aging out endstation entries in the switch Forwarding Database (FDB):

Reduce the number of topology changes by disabling STP on those systems that do not use redundant paths.

Specify that the endstation entries are static or permanent.

Index

A

access levels 2-8
accounts, creating 2-10
`admin` account 2-9
aging entries 5-1
alarm actions 11-19
Alarms, RMON 11-18
area 0, OSPF 9-6
areas, OSPF 9-5
autonegotiation 3-4

B

backbone area, OSPF 9-6
BlackDiamond
- boot option commands (table) 13-5
- components 1-1
- configuration example 1-6
- factory defaults 1-8
- features 1-3
- home page 2-17, 12-3
- image upgrade 13-1
- logging 11-11
- monitoring 11-1
- rebooting 13-2
- resetting to factory defaults 13-3
- RMON features 11-17
- saving configuration changes 13-3
- slot configuration 3-1
- uploading configuration 13-4

blackhole entries 5-2
boot option commands (table) 13-5
BOOTP relay, configuring 8-10
BOOTP, using 2-13
bridge priority 6-5
browser
- controls 12-4
- fonts 12-2
- setting up 12-2

C

command
- history 2-6
- shortcuts 2-3
- syntax, understanding 2-2

common commands (table) 2-6
community strings 2-19
components 1-1
configuration
- primary and secondary 13-3
- saving changes 13-3
- uploading to file 13-4

conventions
- notice icons, About This Guide xvi
- text, About This Guide xvi

creating a QoS profile 7-6

D

default
- passwords 2-9
- settings 1-8
- users 2-9

default STP domain 6-2
default VLAN 4-15
deleting a session 2-16
DHCP relay, configuring 8-10
disabling a port 3-3

disabling route advertising (RIP) 9-4
disabling Telnet 2-16
disconnecting a Telnet session 2-15
Distance Vector Multicast Routing Protocol. *See* DVMRP
distance-vector protocol, description 9-2
domains, Spanning Tree Protocol 6-1
DVMRP
  configuring 10-3
  description 10-2
dynamic entries 5-1
dynamic routes 8-3

E

enabling a port 3-3
errors, port 11-10
ESRP
  active router 8-11
  configuration commands (table) 8-11
  configuring 8-11
  description 8-11
  example 8-12
  groups 8-11
Events, RMON 11-18
Extreme Standby Router Protocol. *See* ESRP
ExtremeWare Vista
  accessing 12-3
  browser controls 12-4
  browser setup 12-2
  capturing screen output 12-7
  description 12-1
  disabling 2-18, 12-1
  enabling 12-1
  fonts 12-2
  home page 2-17, 12-3
  navigating 12-3
  saving changes 12-5
  screen layout 12-3
  screen resolution 12-2
  status messages 12-5
  VLAN configuration 12-2

F

FDB
  adding an entry 5-3
  agingin entries 5-1
  blackhole entries 5-2
  clear and delete commands (table) 5-6
  configuration commands (table) 5-3
  configuring 5-3
  contents 5-1
  creating a permanent entry example 5-4
  displaying 5-5
  dynamic entries 5-1
  entries 5-1
  non-aging entries 5-2
  permanent entries 5-2
  QoS profile association 5-3
  removing entries 5-6
fonts, browser 12-2
forward delay 6-5
Forwarding Database. *See* FDB
full-duplex 1-4

G

GARP VLAN Registration Protocol. *See* GVRP
GVRP
  configuration commands (table) 4-11
  description 4-9
  example 4-9

H

hello time 6-5
`history` command 2-6
History, RMON 11-17
home page 2-17, 12-3
host configuration commands (table) 2-16

I

ICMP configuration commands (table) 8-15
IEEE 802.1Q 4-6
IGMP
  configuration commands (table) 10-4
  description 10-2
  snooping 10-2
image
  downloading 13-1
  primary and secondary 13-2
  upgrading 13-1
interfaces, router 8-2
Internet Group Management Protocol. *See* IGMP
IP address, entering 2-14
IP multicast routing
  configuration commands (table) 10-3
  configuring 10-2
  description 1-6, 10-1
  disabling 10-8
  DVMRP
    configuring 10-3
    description 10-2
  example 10-5
  IGMP
    configuration commands (table) 10-4
    description 10-2
    snooping 10-2
  PIM-DM
    configuring 10-3
  reset and disable commands (table) 10-8

resetting 10-8
settings, displaying 10-7
show commands (table) 10-7
IP multinetting
description 8-5
example 8-7
IP unicast routing
BOOTP relay 8-10
configuration examples 8-16
configuring 8-9
default gateway 8-1
description 1-6
DHCP relay 8-10
disabling 8-20
enabling 8-9
multinetting, description 8-5
multinetting, example 8-7
proxy ARP 8-4
reset and disable commands (table) 8-20
resetting 8-20
router interfaces 8-2
router show commands (table) 8-18
routing table
configuration commands (table) 8-14
dynamic routes 8-3
multiple routes 8-4
populating 8-3
static routes 8-3
settings, displaying 8-18
verifying the configuration 8-9

K

keys
line-editing 2-5
port monitoring 11-11

L

line-editing keys 2-5
link-state database 9-5
link-state protocol, description 9-2
load sharing
configuring 3-7
description 3-7
example 3-9
group combinations (table) 3-8
load-sharing group, description 3-7
master port 3-8
verifying the configuration 3-9
local logging 11-13
log display 11-13
logging
and Telnet 11-13
commands (table) 11-14
description 11-11
fault level 11-12
local 11-13
message 11-12
real-time display 11-13
remote 11-14
subsystem 11-12
timestamp 11-11
logging in 2-9

M

management access 2-8
master port, load sharing 3-8
max age 6-5
MIBs 2-19
mode, QoS 7-2
monitoring the switch 11-1
multiple routes 8-4

N

names, VLANs 4-14
non-aging entries 5-2

O

Open Shortest Path First. *See* OSPF
OSPF
advantages 9-2
area 0 9-6
areas 9-5
backbone area 9-6
configuration commands (table) 9-14
configuration example 9-17
description 9-2, 9-5
disabling 9-20
enabling 8-9
link-state database 9-5
reset and disable commands (table) 9-20
resetting 9-20
router types 9-6
settings, displaying 9-20
show commands (table) 9-20
stub area 9-6
virtual link 9-7

P

PACE 7-5
passwords
default 2-9
forgetting 2-10
path cost 6-6
permanent entries 5-2
PIM-DM
configuration 10-3
`ping` command 2-22

poison reverse 9-3
port
autonegotiation 3-4
commands (table) 3-5
configuring 3-1
enabling and disabling 3-3
errors,viewing 11-10
master port 3-8
monitoring display keys 11-11
priority, STP 6-6
receive errors 11-10
statistics, viewing 11-9
STP state, displaying 6-9
STPD membership 6-2
transmit errors 11-10
Port Queue Monitor. *See* PQM
port-based VLANs 4-2
port-mirroring
configuration commands (table) 3-10
description 3-9
example 3-10
virtual port 3-10
PQM
commands (table) 7-8
description 7-7
primary image 13-2
profiles, QoS 7-2
protocol filters 4-12
protocol-based VLANs 4-11
proxy ARP, description 8-4

Q

QoS
building blocks 7-1
configuration commands (table) 7-8
configuration examples 7-9
configuring 7-8
default QoS profiles 7-2
description 1-5, 7-1
FDB entry association 5-3
information, displaying 7-10
mode 7-1
PACE recognition 7-5
Port Queue Monitor (PQM), description 7-7
precedence 7-5
prioritization 7-6
profiles
configuring 7-9
creating 7-6
deleting 7-11
description 7-2
resetting 7-11
traffic classification 7-1
traffic groupings, description 7-3
Quality of Service. *See* QoS

R

rebooting 13-2
receive errors 11-10
remote logging 11-14
Remote Monitoring. *See* RMON
reset to factory defaults 13-3
RIP
advantages 9-2
configuration commands (table) 9-9
configuration example 9-11
description 9-2, 9-3
disabling route advertising 9-4
enabling 8-9
limitations 9-2
poison reverse 9-3
reset and disable commands (table) 9-14
routing table entries 9-3
settings, displaying 9-13
show commands (table) 9-13
split horizon 9-3
triggered updates 9-4
version 2 9-4
RMON
alarm actions 11-19
Alarms group 11-18
Events group 11-18
features supported 11-17
History group 11-17
probe 11-17
Statistics group 11-17
router interfaces 8-2
router types, OSPF 9-6
Routing Information Protocol. *See* RIP
routing table, populating 8-3
routing. *See* IP unicast routing

S

saving changes using ExtremeWare Vista 12-5
saving configuration changes 13-3
screen resolution, ExtremeWare Vista 12-2
secondary image 13-2
sessions, deleting 2-16
shortcuts, command 2-3
show commands 11-1
Simple Network Management Protocol. *See* SNMP
slot
automatic configuration 3-1
clearing 3-2
manually configuring 3-2
mismatch 3-2
SNAP protocol 4-13
SNMP
authorized managers 2-19
community strings 2-19
configuration commands (table) 2-20

configuring 2-19
reset and disable commands (table) 2-22
settings, displaying 2-21
supported MIBs 2-19
trap receivers 2-19
using 2-18
Spanning Tree Protocol. *See* STP
speed, ports 3-4
split horizon 9-3
static routes 8-3
statistics, port 11-9
Statistics, RMON 11-17
status monitoring 11-1
STP
and VLANs 6-2
configurable parameters 6-5
configuration commands (table) 6-6
configuration example 6-8
configuring 6-5
default domain 6-2
description 1-5
disable and reset commands (table) 6-9
displaying settings 6-8
domains 6-1
examples 6-2
overview 6-1
port state, displaying 6-9
stub area, OSPF 9-6
switch logging 11-11
switch monitoring commands (table) 11-1
syntax, understanding 2-2
syslog host 11-14

T

tagging, VLAN 4-6
Telnet
disabling 2-16
disconnecting a session 2-15
logging 11-13
using 2-12
TFTP
server 13-1
using 13-4
`traceroute` command 2-23
traffic groupings, QoS 7-3
transmit errors 11-10
triggered updates 9-4
trunks 4-6

U

upgrading the image 13-1
uploading the configuration 13-4
users
access levels 2-8
creating 2-10
default 2-9
viewing 2-10

V

viewing accounts 2-10
Virtual LANs. *See* VLANs
virtual link, OSPF 9-7
VLAN tagging 4-6
VLANs
and ExtremeWare Vista 12-2
and STP 6-2
assigning a tag 4-6
benefits 4-1
configuration commands (table) 4-15
configuration examples 4-17
configuring 4-15
default 4-15
delete and reset commands (table) 4-20
description 1-4
disabling route advertising 9-4
displaying settings 4-18
mixing port-based and tagged 4-9
names 4-14
port-based 4-2
protocol filters 4-12
protocol-based 4-11
restoring default values 4-20
routing 8-9
tagged 4-6
trunks 4-6
types 4-2

W

Web access
disabling 2-18, 12-1
enabling 12-1

X

xmodem 13-2

Index of Commands

C

clear counters 11-16
clear fdb 5-6
clear igmp snooping 10-8
clear iparp 2-16, 8-13, 8-20
clear ipfdb 8-14, 8-20
clear ipmc cache 10-8
clear log 11-16
clear session 2-7, 2-16
clear slot 3-2
clear slot module 3-5
config account 2-6
config banner 2-6
config bootprelay add 8-13
config bootprelay delete 8-13
config dot1q ethertype 4-15
config dvmrp add 10-3
config dvmrp delete 10-3
config dvmrp timer 10-4
config dvmrp vlan 10-3
config fdb agingtime 5-4
config gvrp 4-11
config igmp query_interval 10-4
config igmp snooping 10-5
config iparp add 2-16, 8-13
config iparp add proxy 8-13
config iparp delete 2-16, 8-13
config iparp delete proxy 8-13
config ipmc cache timeout 10-4
config ipqos add 8-14
config ipqos delete 8-14
config iproute add 2-17, 8-14
config iproute add blackhole 8-14
config iproute add default 2-17, 8-15
config iproute delete 2-17, 8-14
config iproute delete blackhole 8-14
config iproute delete default 2-17, 8-15
config irdp 8-15, 8-16
config log display 11-15
config mirroring add 3-10
config mirroring delete 3-10
config ospf add virtual-link 9-15
config ospf add vlan 9-15
config ospf area add range 9-15
config ospf area delete range 9-16
config ospf area normal 9-15
config ospf area stub 9-15
config ospf authentication 9-15
config ospf cost 9-15
config ospf delete virtual-link 9-15
config ospf delete vlan 9-15, 9-20
config ospf priority 9-15
config ospf routerid 9-16
config ospf timer 9-16
config ospf vlan 9-15
config port 2-6
config port auto off 3-5
config port auto on 3-5
config port esrp group 8-14
config port esrp priority 8-14
config port esrp timer 8-12, 8-14
config port qosprofile 3-5, 7-9
config protocol 4-16

config qosmode 7-8
config qosprofile 7-9
config rip add 9-10
config rip delete 9-10, 9-14
config rip garbagetime 9-10
config rip routetimeout 9-10
config rip rxmode 9-10
config rip txmode 9-11
config rip updatetime 9-11
config slot 3-2
config slot module 2-6, 3-5
config snmp add 2-20
config snmp add trapreceiver 2-20
config snmp delete 2-20
config snmp delete trapreceiver 2-20
config snmp syscontact 2-21
config snmp syslocation 2-21
config snmp sysname 2-21
config stpd add vlan 6-6
config stpd forwarddelay 6-6
config stpd hellotime 6-6
config stpd maxage 6-7
config stpd port cost 6-7
config stpd port priority 6-7
config stpd priority 6-7
config syslog 11-15
config time 2-6
config vlan add port 4-16
config vlan delete port 4-16
config vlan esrp group 8-11
config vlan esrp priority 8-11
config vlan ipaddress 2-7, 4-16
config vlan protocol 4-16
config vlan qosprofile 4-16, 7-9
config vlan tag 4-16
copy configuration 13-6
copy image 13-5
create account 2-6
create fdbentry 5-3
create ospf area 9-14
create protocol 4-15
create qosprofile 7-8
create stpd 6-6
create vlan 2-6, 4-15

D

delete account 2-7
delete fdbentry 5-6
delete ospf area 9-20
delete protocol 4-20
delete qosprofile 7-11
delete stpd 6-9
delete vlan 2-7, 4-20
disable bootp 2-7, 8-13, 8-20
disable bootprelay 8-13, 8-20
disable dvmrp 10-8
disable esrp port 8-12, 8-14
disable gvrp 4-11
disable icmp redirects 8-16, 8-20
disable icmp unreachables 8-16, 8-20
disable icmp useredirects 8-16, 8-20
disable idletimeout 2-7
disable igmp 10-8
disable igmp snooping 10-8
disable ignore-stp 4-20
disable ipforwarding 8-13, 8-20
disable ipforwarding broadcast 8-13, 8-20
disable ipmcforwarding 10-8
disable iproute sharing 8-15
disable irdp 8-16, 8-20
disable learning port 3-5, 5-4
disable log display 11-15
disable mirroring 3-10
disable multinetting 8-13
disable ospf 9-20
disable ospf export rip 9-20
disable ospf export static 9-20
disable pace 7-5, 7-9
disable port 2-7, 3-3, 3-5
disable qosmonitor 7-8
disable rip 9-14
disable rip aggregation 9-14
disable rip export ospf 9-14
disable rip export static 9-14
disable rip poisonreverse 9-14
disable rip splithorizon 9-14
disable rip triggerupdate 9-14
disable sharing 3-5
disable snmp access 2-22
disable snmp trap 2-22
disable stpd 6-9
disable stpd port 6-9
disable syslog 11-15
disable telnet 2-7
disable web 2-7, 12-1
download bootrom 13-5
download config 13-5
download image 13-5

E

enable bootp 2-7, 8-12
enable bootprelay 8-12
enable dvmrp 10-3
enable esrp port 8-14
enable esrp vlan 8-11
enable gvrp 4-11
enable icmp redirects 8-15
enable icmp unreachables 8-15
enable icmp useredirects 8-15
enable idletimeout 2-7
enable igmp 10-4
enable ignore-stp 4-15
enable ipforwarding 8-12
enable ipforwarding broadcast 8-13
enable ipmcforwarding 10-3
enable iproute sharing 8-14
enable irdp 8-15
enable learning port 3-5, 5-4
enable log display 11-15
enable mirroring 3-10
enable multinetting 8-13
enable ospf 9-14
enable ospf export rip 9-9, 9-15
enable ospf export static type 9-14
enable pace 7-5, 7-8
enable port 3-3, 3-5
enable qosmonitor 7-8
enable rip 9-9
enable rip aggregation 9-9
enable rip export ospf 9-9, 9-10
enable rip export static 9-10
enable rip poisonreverse 9-10
enable rip splithorizon 9-10
enable rip triggerupdate 9-10
enable sharing 3-5
enable snmp access 2-20
enable snmp trap 2-20
enable stpd 6-6
enable stpd port 6-6
enable syslog 11-15
enable telnet 2-16
enable web 2-18, 12-1

H

history 2-6

L

logout 2-15

P

ping 2-22

Q

quit 2-15

R

reboot 13-2, 13-6
run diag 11-8

S

save config 13-6
show account 2-10, 11-1
show banner 2-7, 11-1
show config 11-2, 13-5
show diag 11-2
show dvmrp 10-7, 11-2
show fdb 5-5, 11-2
show gvrp 4-11, 11-2
show igmp snooping 10-7, 11-2
show iparp 2-17, 8-19, 11-2
show iparp proxy 8-18, 11-2
show ipconfig 2-17, 8-10, 8-19, 11-2
show ipfdb 8-19, 11-2
show ipmc cache 10-7, 11-3
show ipqos 8-19, 11-3
show iproute 8-19, 11-3
show ipstats 2-17, 8-19, 11-3
show log 11-3, 11-16
show log config 11-3, 11-16
show management 2-21, 11-4
show memory 11-4
show mirroring 3-10, 11-4
show ospf 9-20, 11-4
show ospf area 9-20, 11-4
show ospf interfaces 9-20, 11-4
show ospf lsdb 9-20, 11-4
show ospf virtual-link 9-20, 11-4
show port collisions 3-5, 11-4
show port config 3-6, 3-9, 11-4
show port information 3-6, 11-5
show port packet 3-6, 11-5
show port qosmonitor 3-6, 7-8, 11-5
show port rxerrors 3-6, 11-5
show port stats 3-6, 11-5
show port txerrors 3-6, 11-5
show port utilization 3-6, 11-5
show protocol 4-19, 11-5

show qosprofile 7-10, 11-5
show rip 9-13, 11-5
show rip stat 9-13, 11-6
show session 2-15, 11-6
show slot 3-2, 3-7, 11-6
show stpd 6-8, 11-6
show stpd port 6-9, 11-6
show switch 11-7
show version 11-7
show vlan 4-18, 11-7

T

telnet 2-12
traceroute 2-23

U

unconfig dvmrp 10-8
unconfig icmp 8-16, 8-20
unconfig igmp 10-8
unconfig irdp 8-16, 8-20
unconfig management 2-22
unconfig rip 9-14
unconfig stpd 6-9
unconfig switch 2-7
unconfig vlan ipaddress 4-20
upload config 13-6
use config 13-6
use image 13-6